



US009715388B2

(12) **United States Patent**  
**Chung et al.**

(10) **Patent No.:** **US 9,715,388 B2**  
(45) **Date of Patent:** **Jul. 25, 2017**

(54) **INSTRUCTION AND LOGIC TO MONITOR LOOP TRIP COUNT AND REMOVE LOOP OPTIMIZATIONS**

(75) Inventors: **Jaewoong Chung**, Sunnyvale, CA (US); **Hyunchul Park**, Sunnyvale, CA (US); **Hongbo Rong**, Cupertino, CA (US); **Cheng Wang**, San Ramon, CA (US); **Youfeng Wu**, Palo Alto, CA (US)

(73) Assignee: **Intel Corporation**, Santa Clara, CA (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 646 days.

(21) Appl. No.: **13/996,861**

(22) PCT Filed: **Mar. 30, 2012**

(86) PCT No.: **PCT/US2012/031711**  
§ 371 (c)(1), (2), (4) Date: **Jun. 21, 2013**

(87) PCT Pub. No.: **WO2013/147896**  
PCT Pub. Date: **Oct. 3, 2013**

(65) **Prior Publication Data**  
US 2014/0208085 A1 Jul. 24, 2014

(51) **Int. Cl.**  
**G06F 15/00** (2006.01)  
**G06F 7/38** (2006.01)  
**G06F 9/00** (2006.01)  
**G06F 9/44** (2006.01)  
**G06F 9/32** (2006.01)  
**G06F 9/30** (2006.01)  
**G06F 9/38** (2006.01)

**G06F 11/34** (2006.01)  
**G06F 9/45** (2006.01)  
(52) **U.S. Cl.**  
CPC ..... **G06F 9/325** (2013.01); **G06F 9/30072** (2013.01); **G06F 9/3842** (2013.01); **G06F 9/3857** (2013.01); **G06F 11/348** (2013.01); **G06F 11/3409** (2013.01); **G06F 8/443** (2013.01); **G06F 2201/88** (2013.01)

(58) **Field of Classification Search**  
None  
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,909,573 A 6/1999 Sheaffer et al.  
7,318,223 B2 1/2008 Blainey et al.  
(Continued)

OTHER PUBLICATIONS

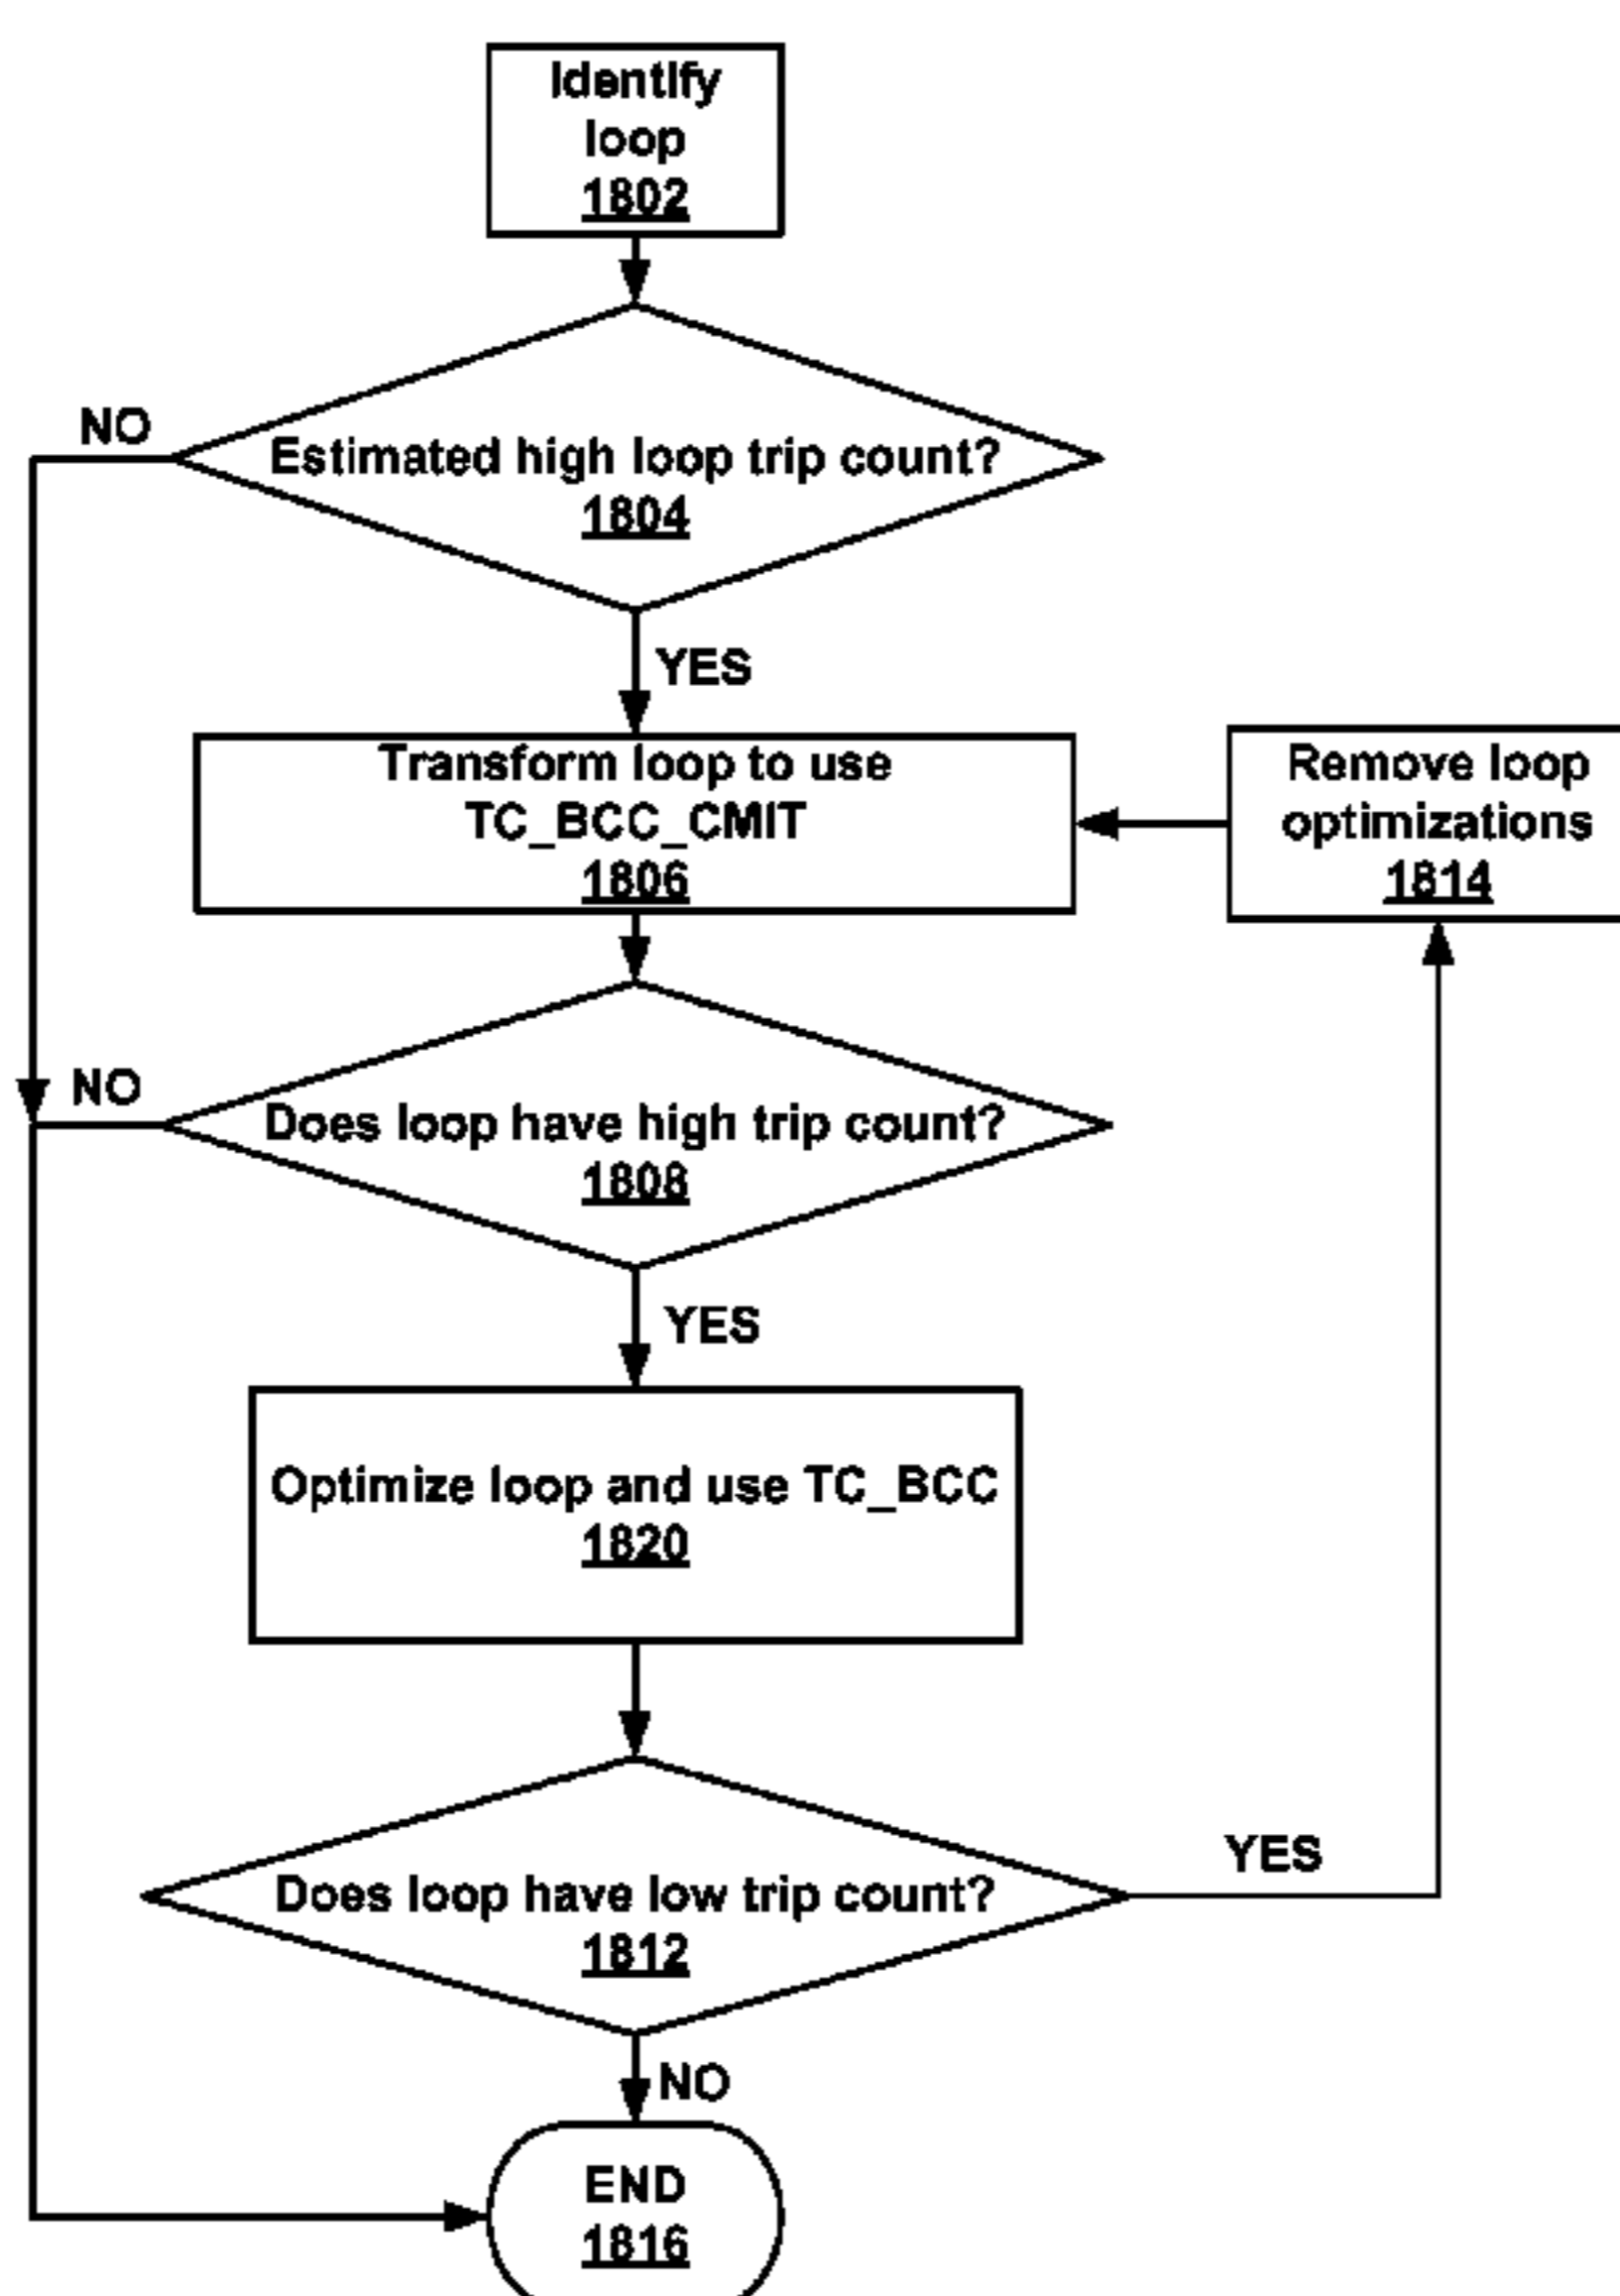
Shukla (Lightweight, Cross-Procedure Tracing for Runtime Optimization, 2003, 58 pages).  
(Continued)

*Primary Examiner* — George Giroux  
(74) *Attorney, Agent, or Firm* — Patent Capital Group

(57) **ABSTRACT**

Logic and instruction to monitor loop trip count are disclosed. Loop trip count information of a loop may be stored in a dedicated hardware buffer. Average loop trip count of the loop may be calculated based on the stored loop trip count information. Based on the average trip count, loop optimizations may be removed from the loop. The stored loop trip count information may include an identifier identifying the loop, a total loop trip count of the loop, and an exit count of the loop.

**26 Claims, 24 Drawing Sheets**



(56)

**References Cited**

U.S. PATENT DOCUMENTS

2001/0044930 A1 11/2001 Miyata et al.  
2005/0240896 A1 10/2005 Wu et al.  
2007/0169030 A1\* 7/2007 Tarditi, Jr. .... G06F 9/466  
717/140

OTHER PUBLICATIONS

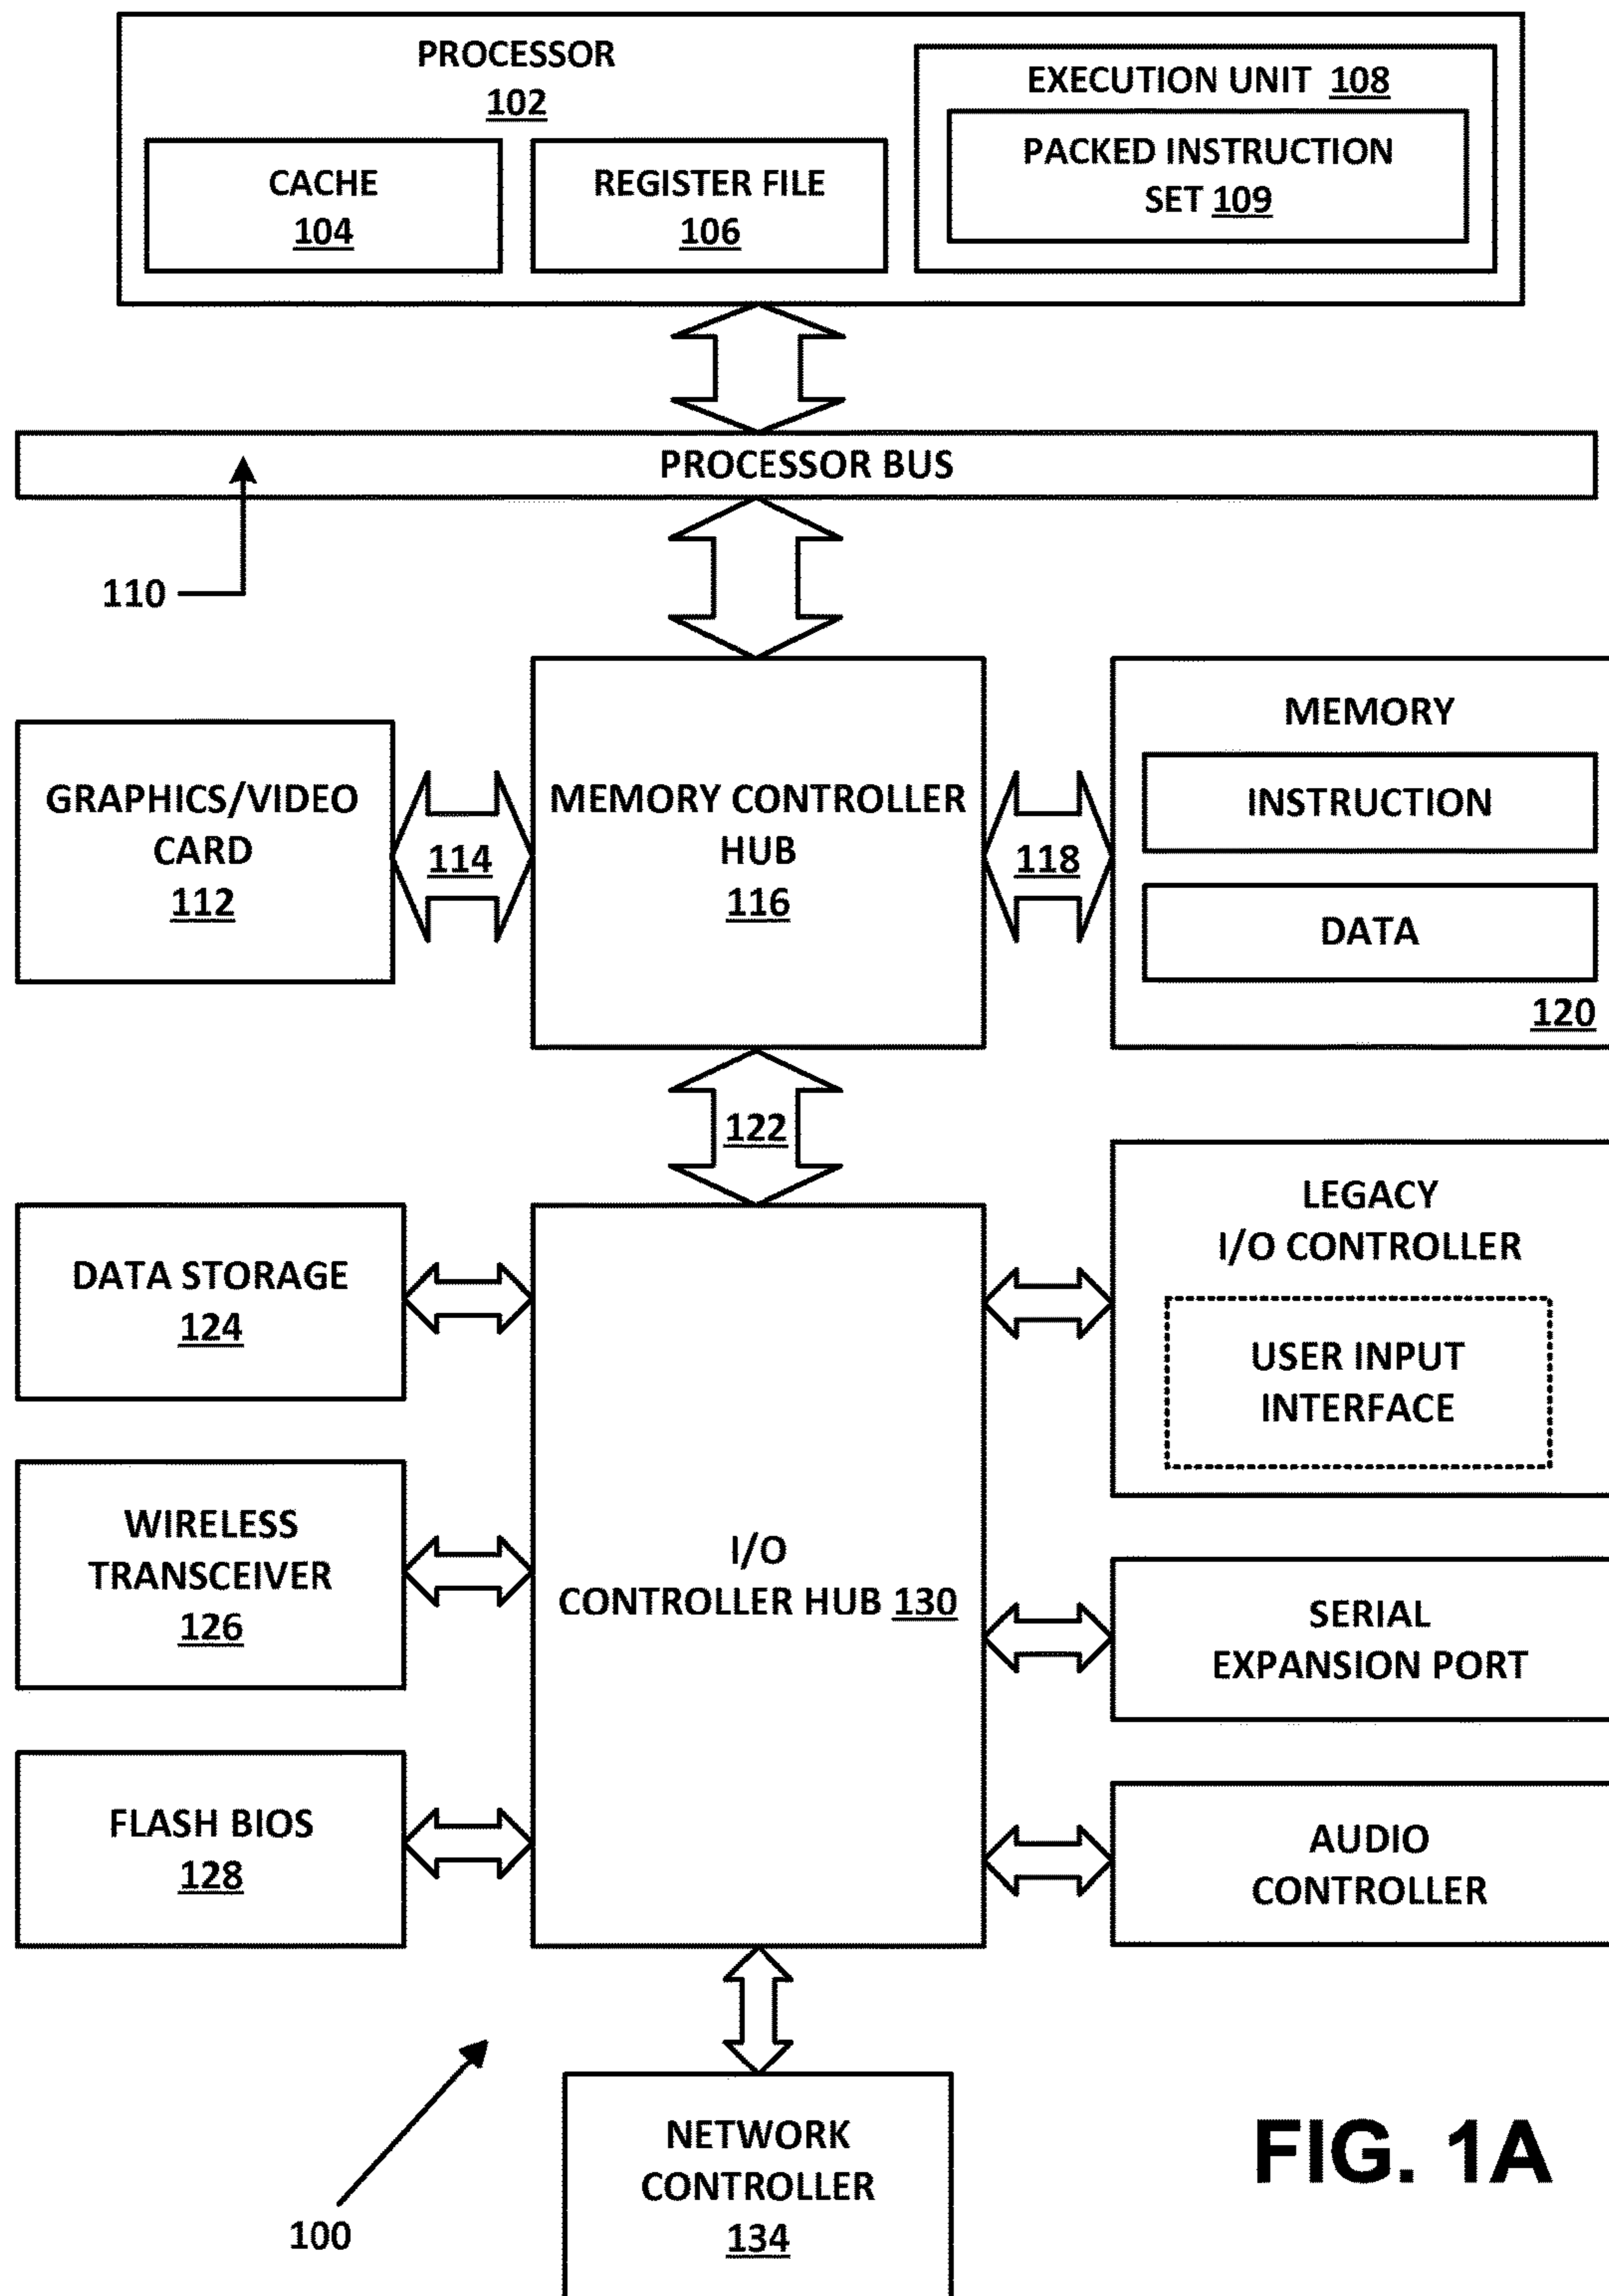
Wang et al. (Y-Branches: When You Come to a Fork in the Road, Take It; Oct. 2003, pp. 1-11).\*

Cavazos et al. (Rapidly Selecting Good Compiler Optimizations using Performance counters, Mar. 2007, pp. 1-13).\*

Zhao et al (Predicting the Impact of Optimizations for Embedded Systems, Jun. 2003, pp. 1-11).\*

International Search Report and Written Opinion dated Nov. 29, 2012 in International Application No. PCT/US2012/031711, filed Mar. 30, 2012.

\* cited by examiner



**FIG. 1A**




FIG. 1B




FIG. 1C




FIG. 2




FIG. 3A






FIG. 3B







UNSIGNED PACKED BYTE REPRESENTATION 344



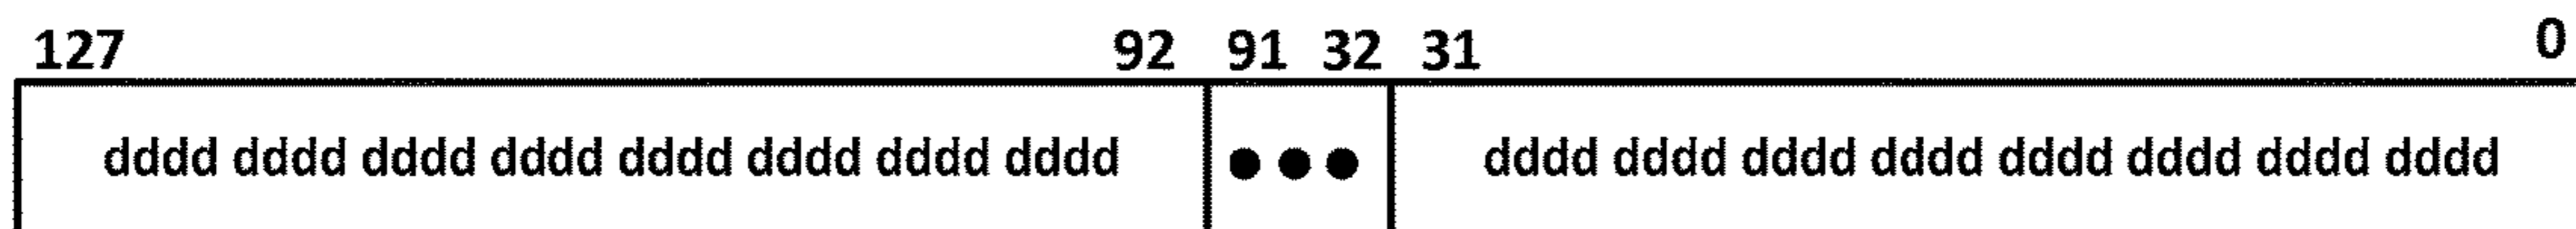
SIGNED PACKED BYTE REPRESENTATION 345



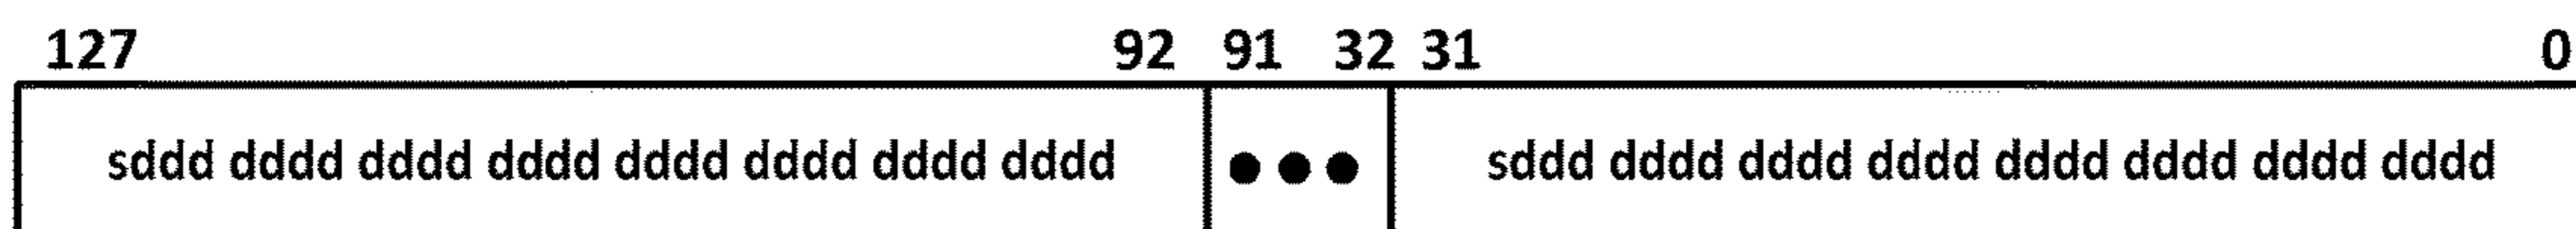
UNSIGNED PACKED WORD REPRESENTATION 346



SIGNED PACKED WORD REPRESENTATION 347



UNSIGNED PACKED DOUBLEWORD REPRESENTATION 348



SIGNED PACKED DOUBLEWORD REPRESENTATION 349

FIG. 3C




FIG. 3D




FIG. 3E


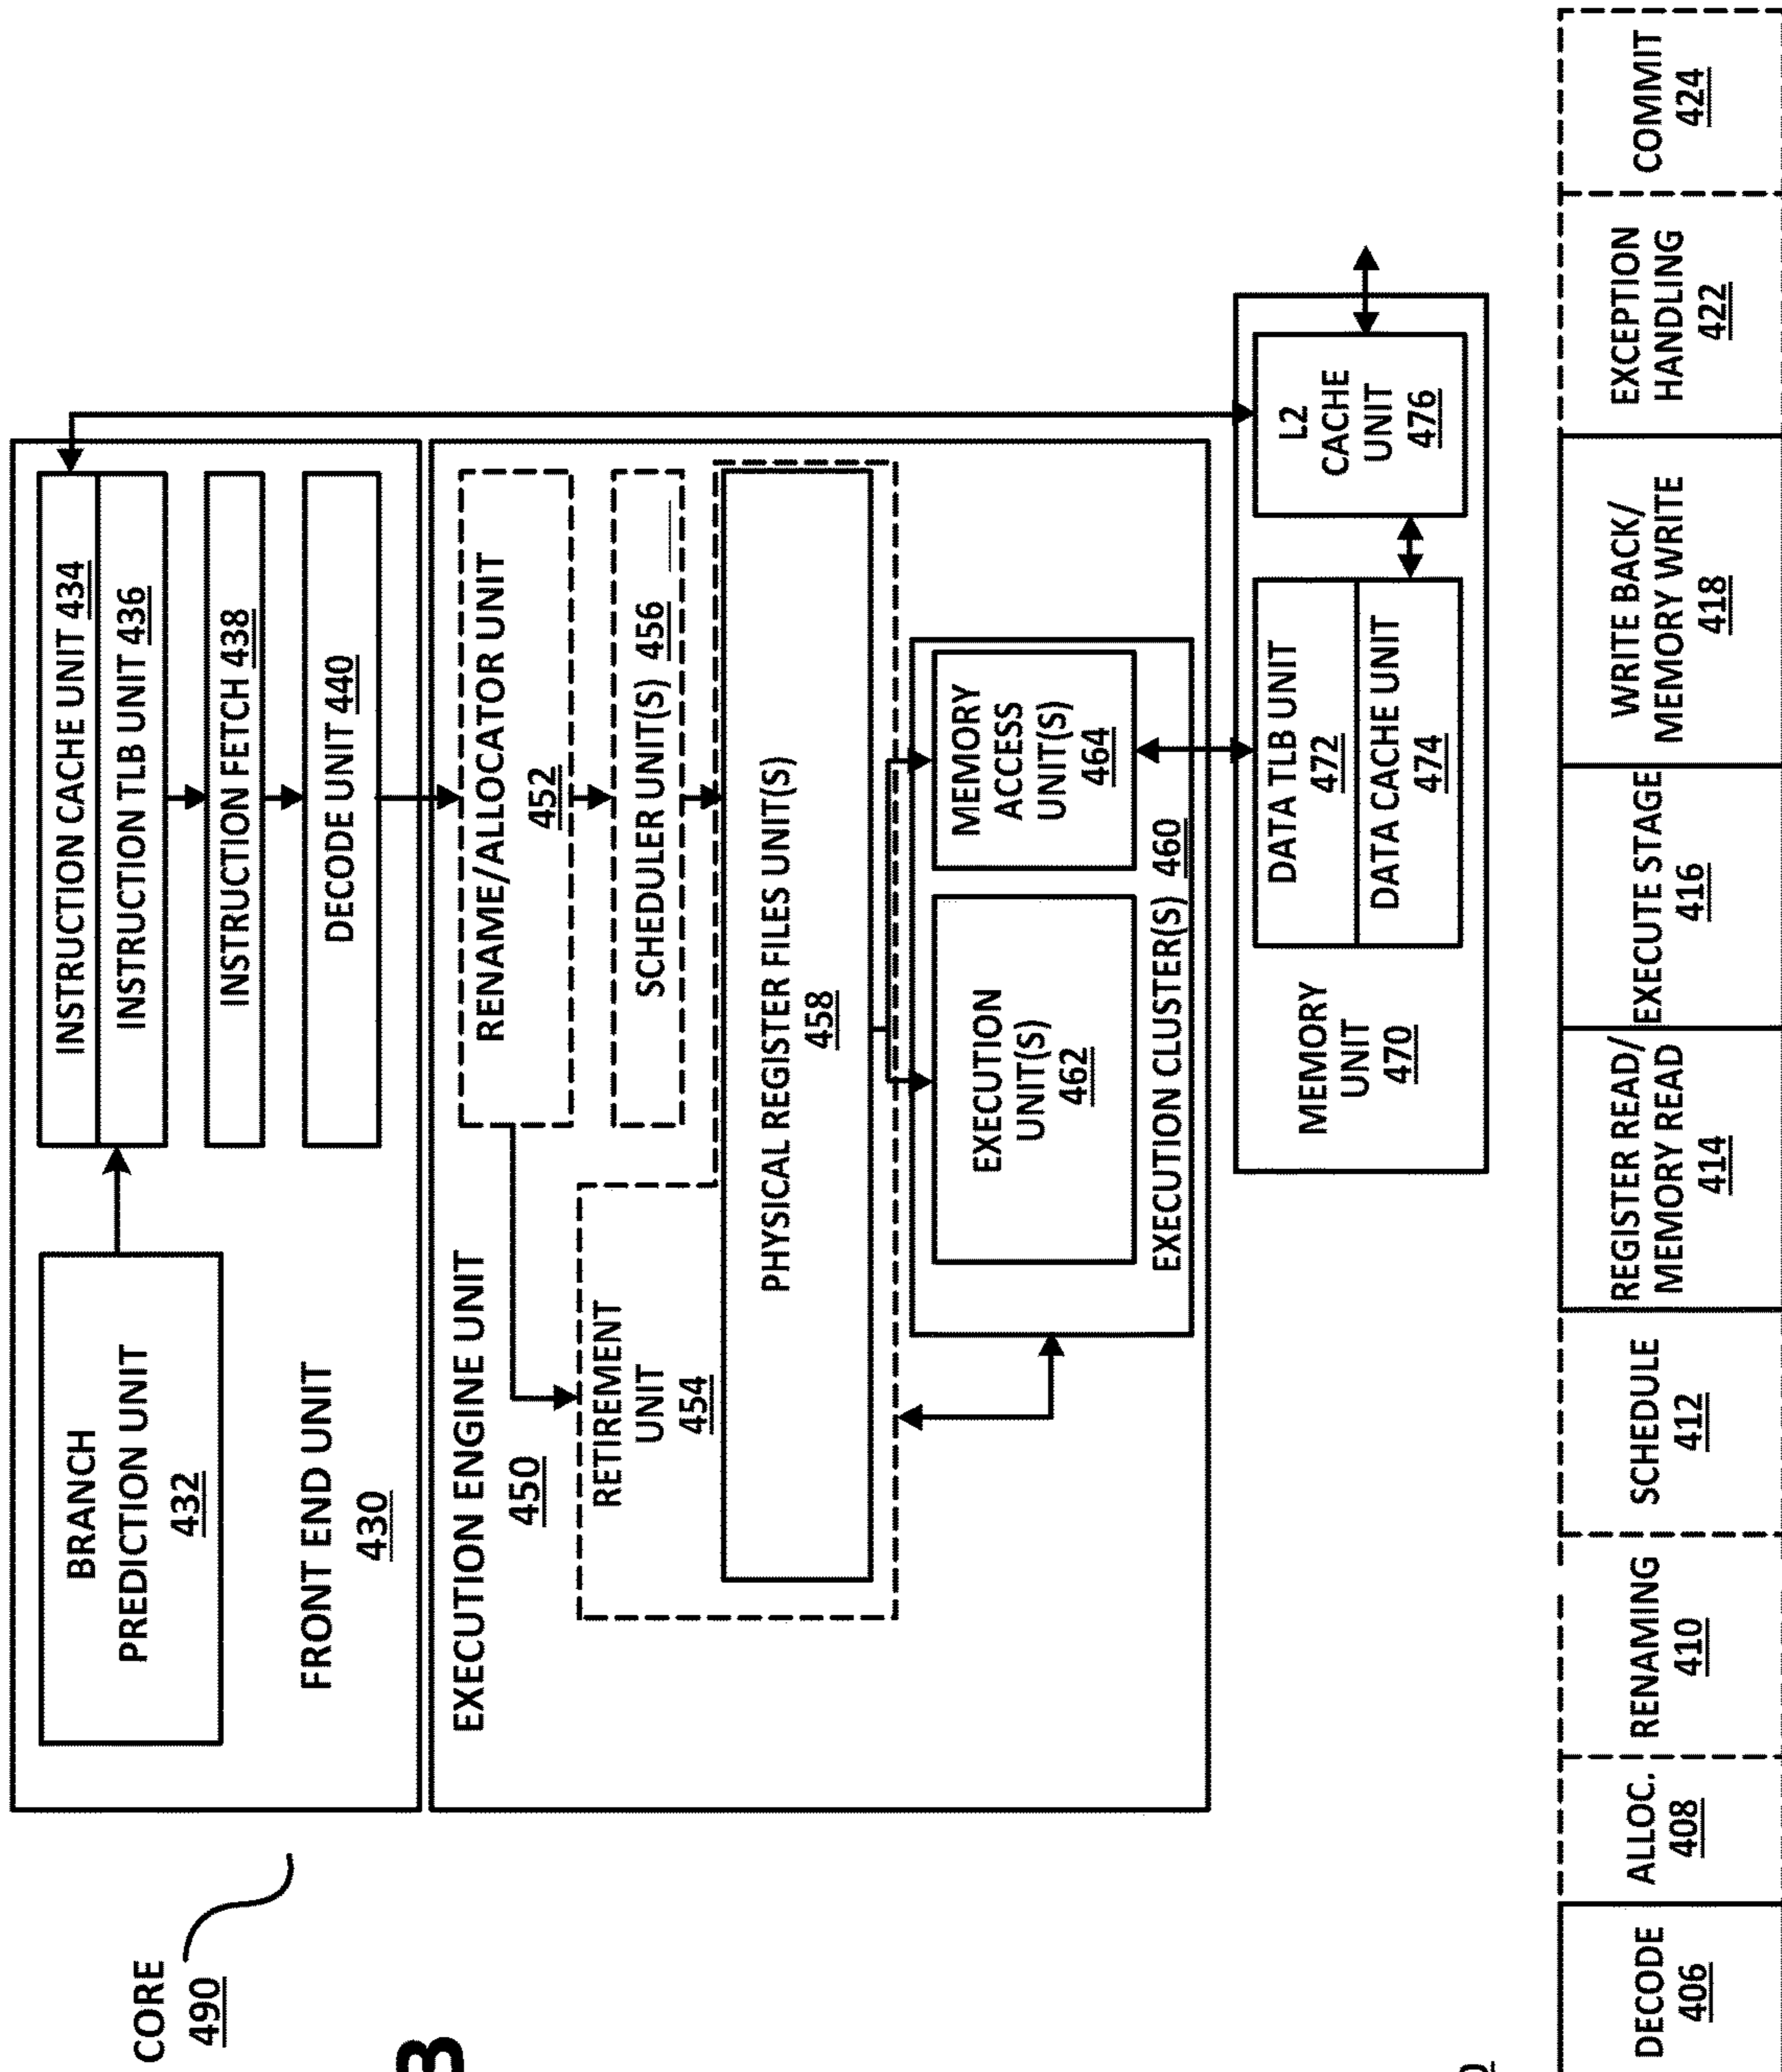


FIG. 3F



**FIG. 4B**

**FIG. 4A**





FIG. 5



**FIG. 6**




FIG. 7





FIG. 8





**FIG. 9**






FIG. 10



**FIG. 11**



**FIG. 12**

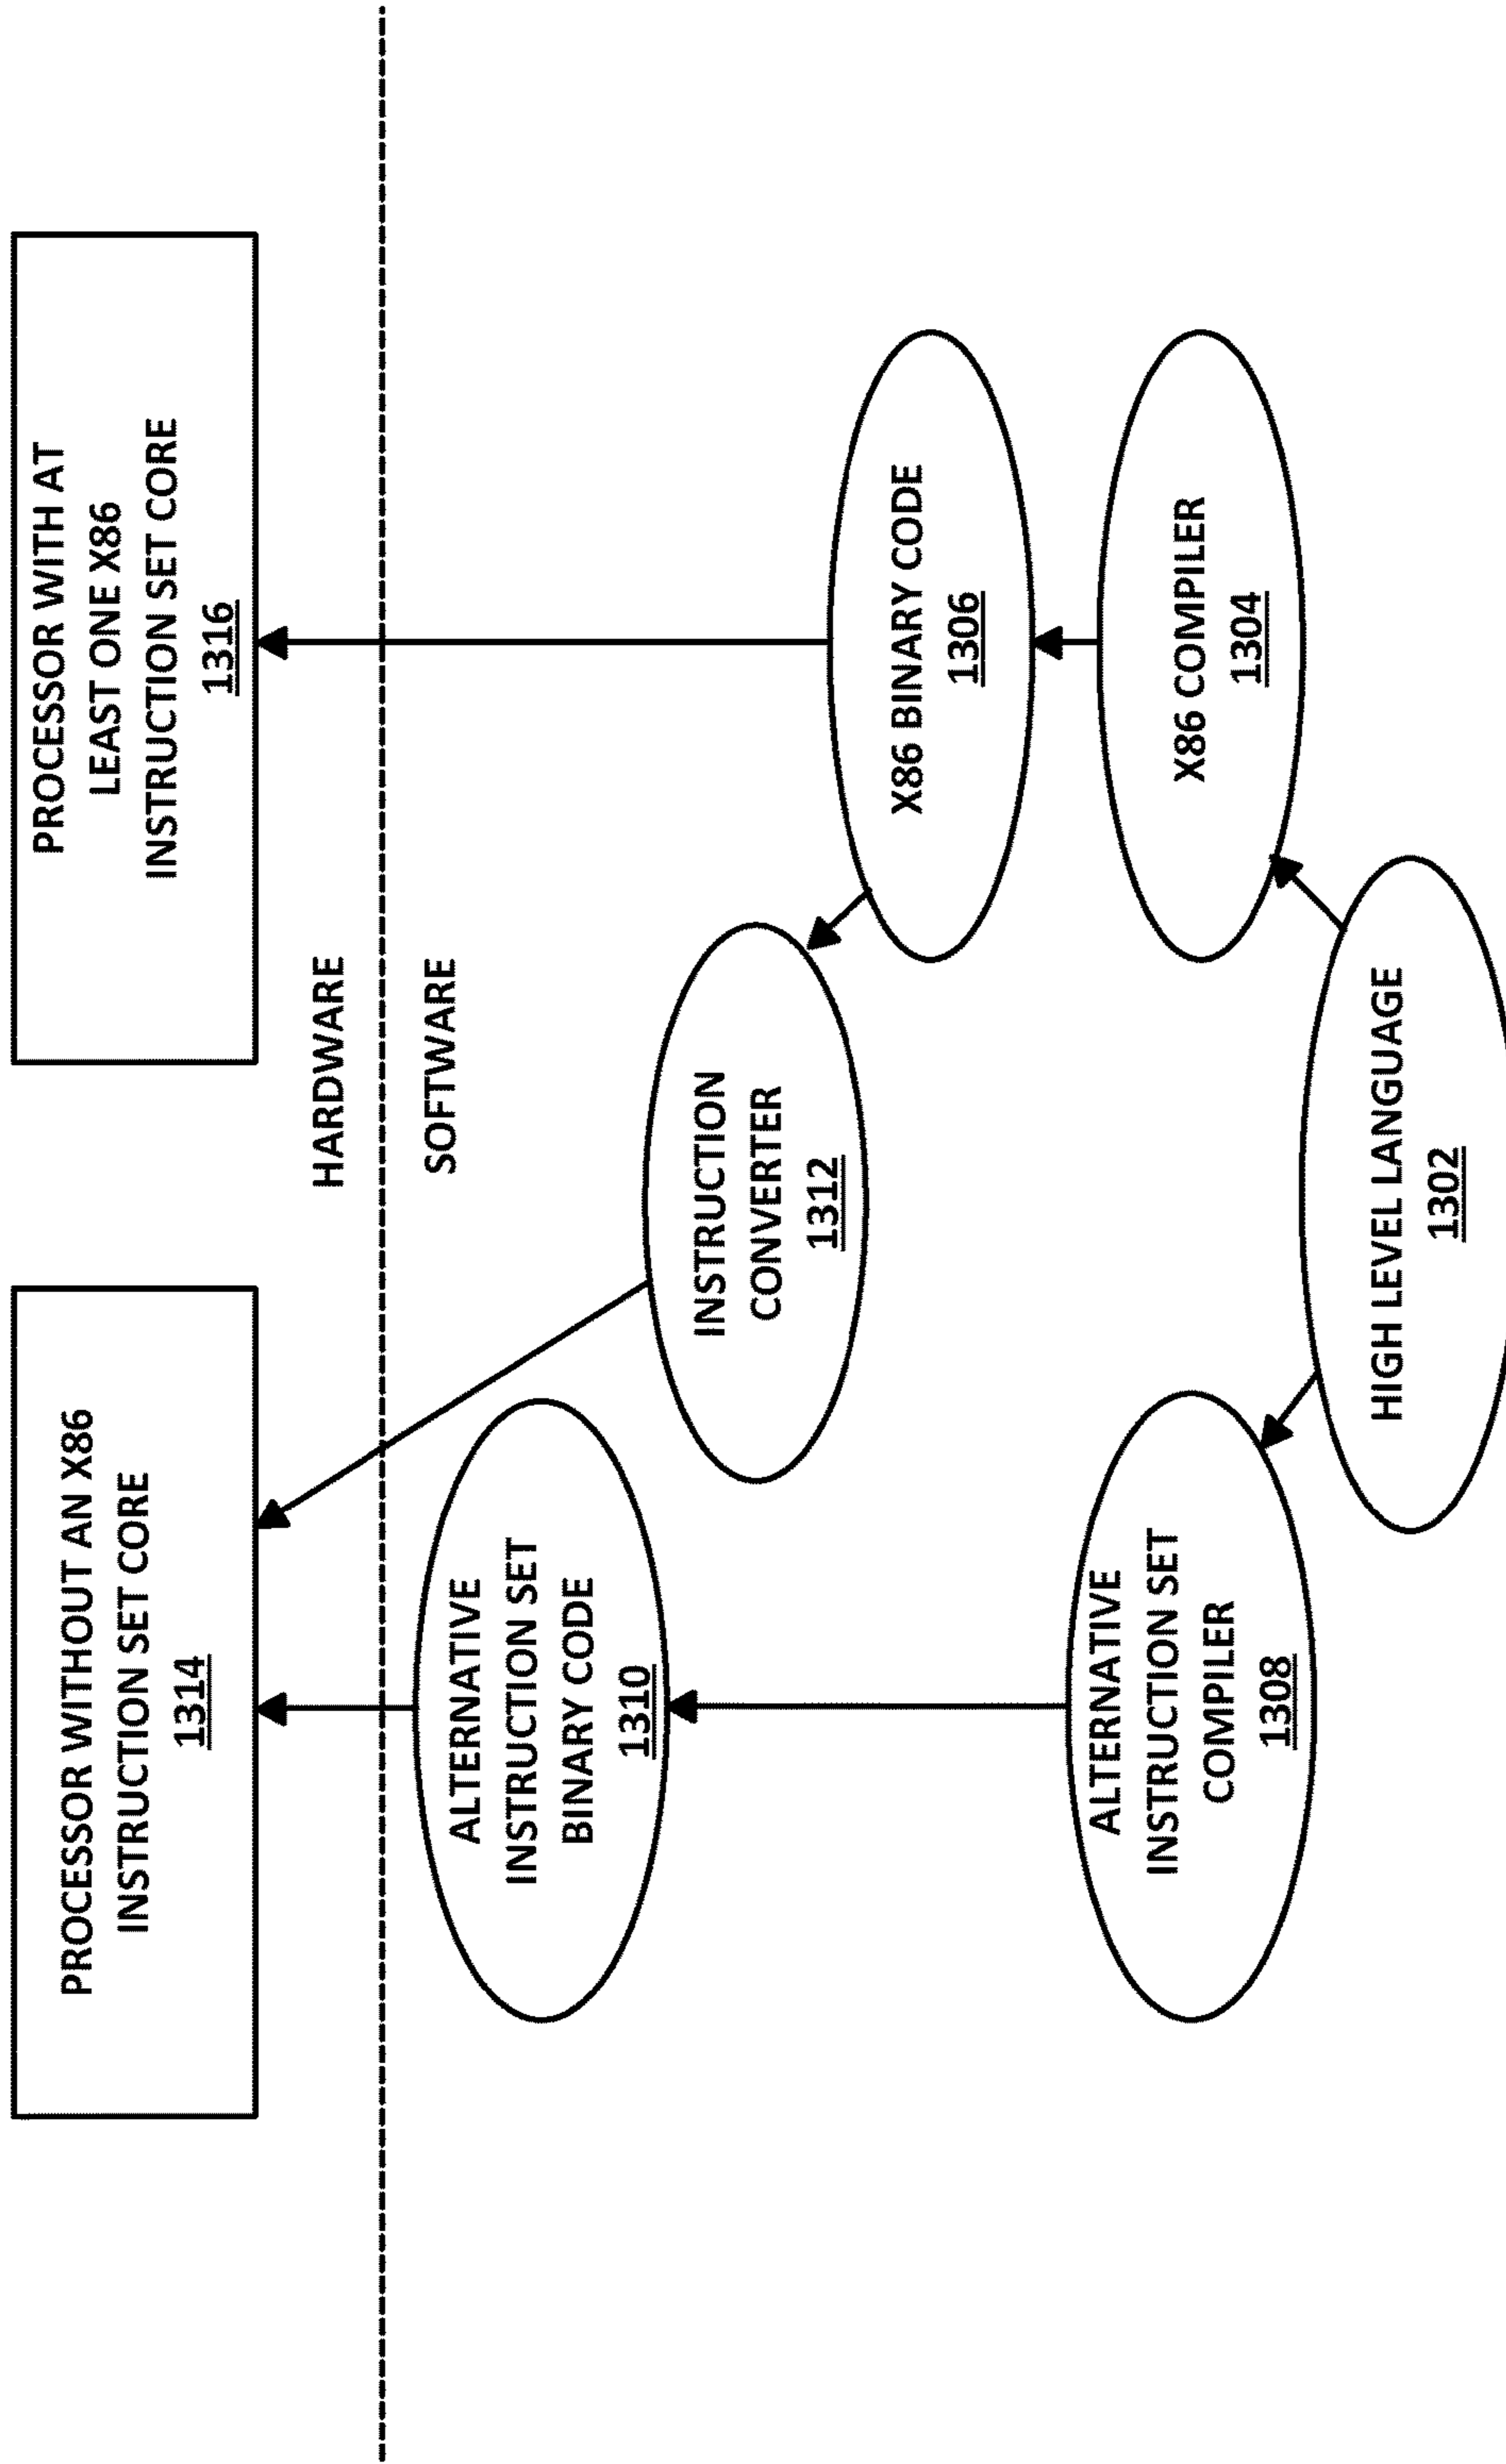


FIG. 13

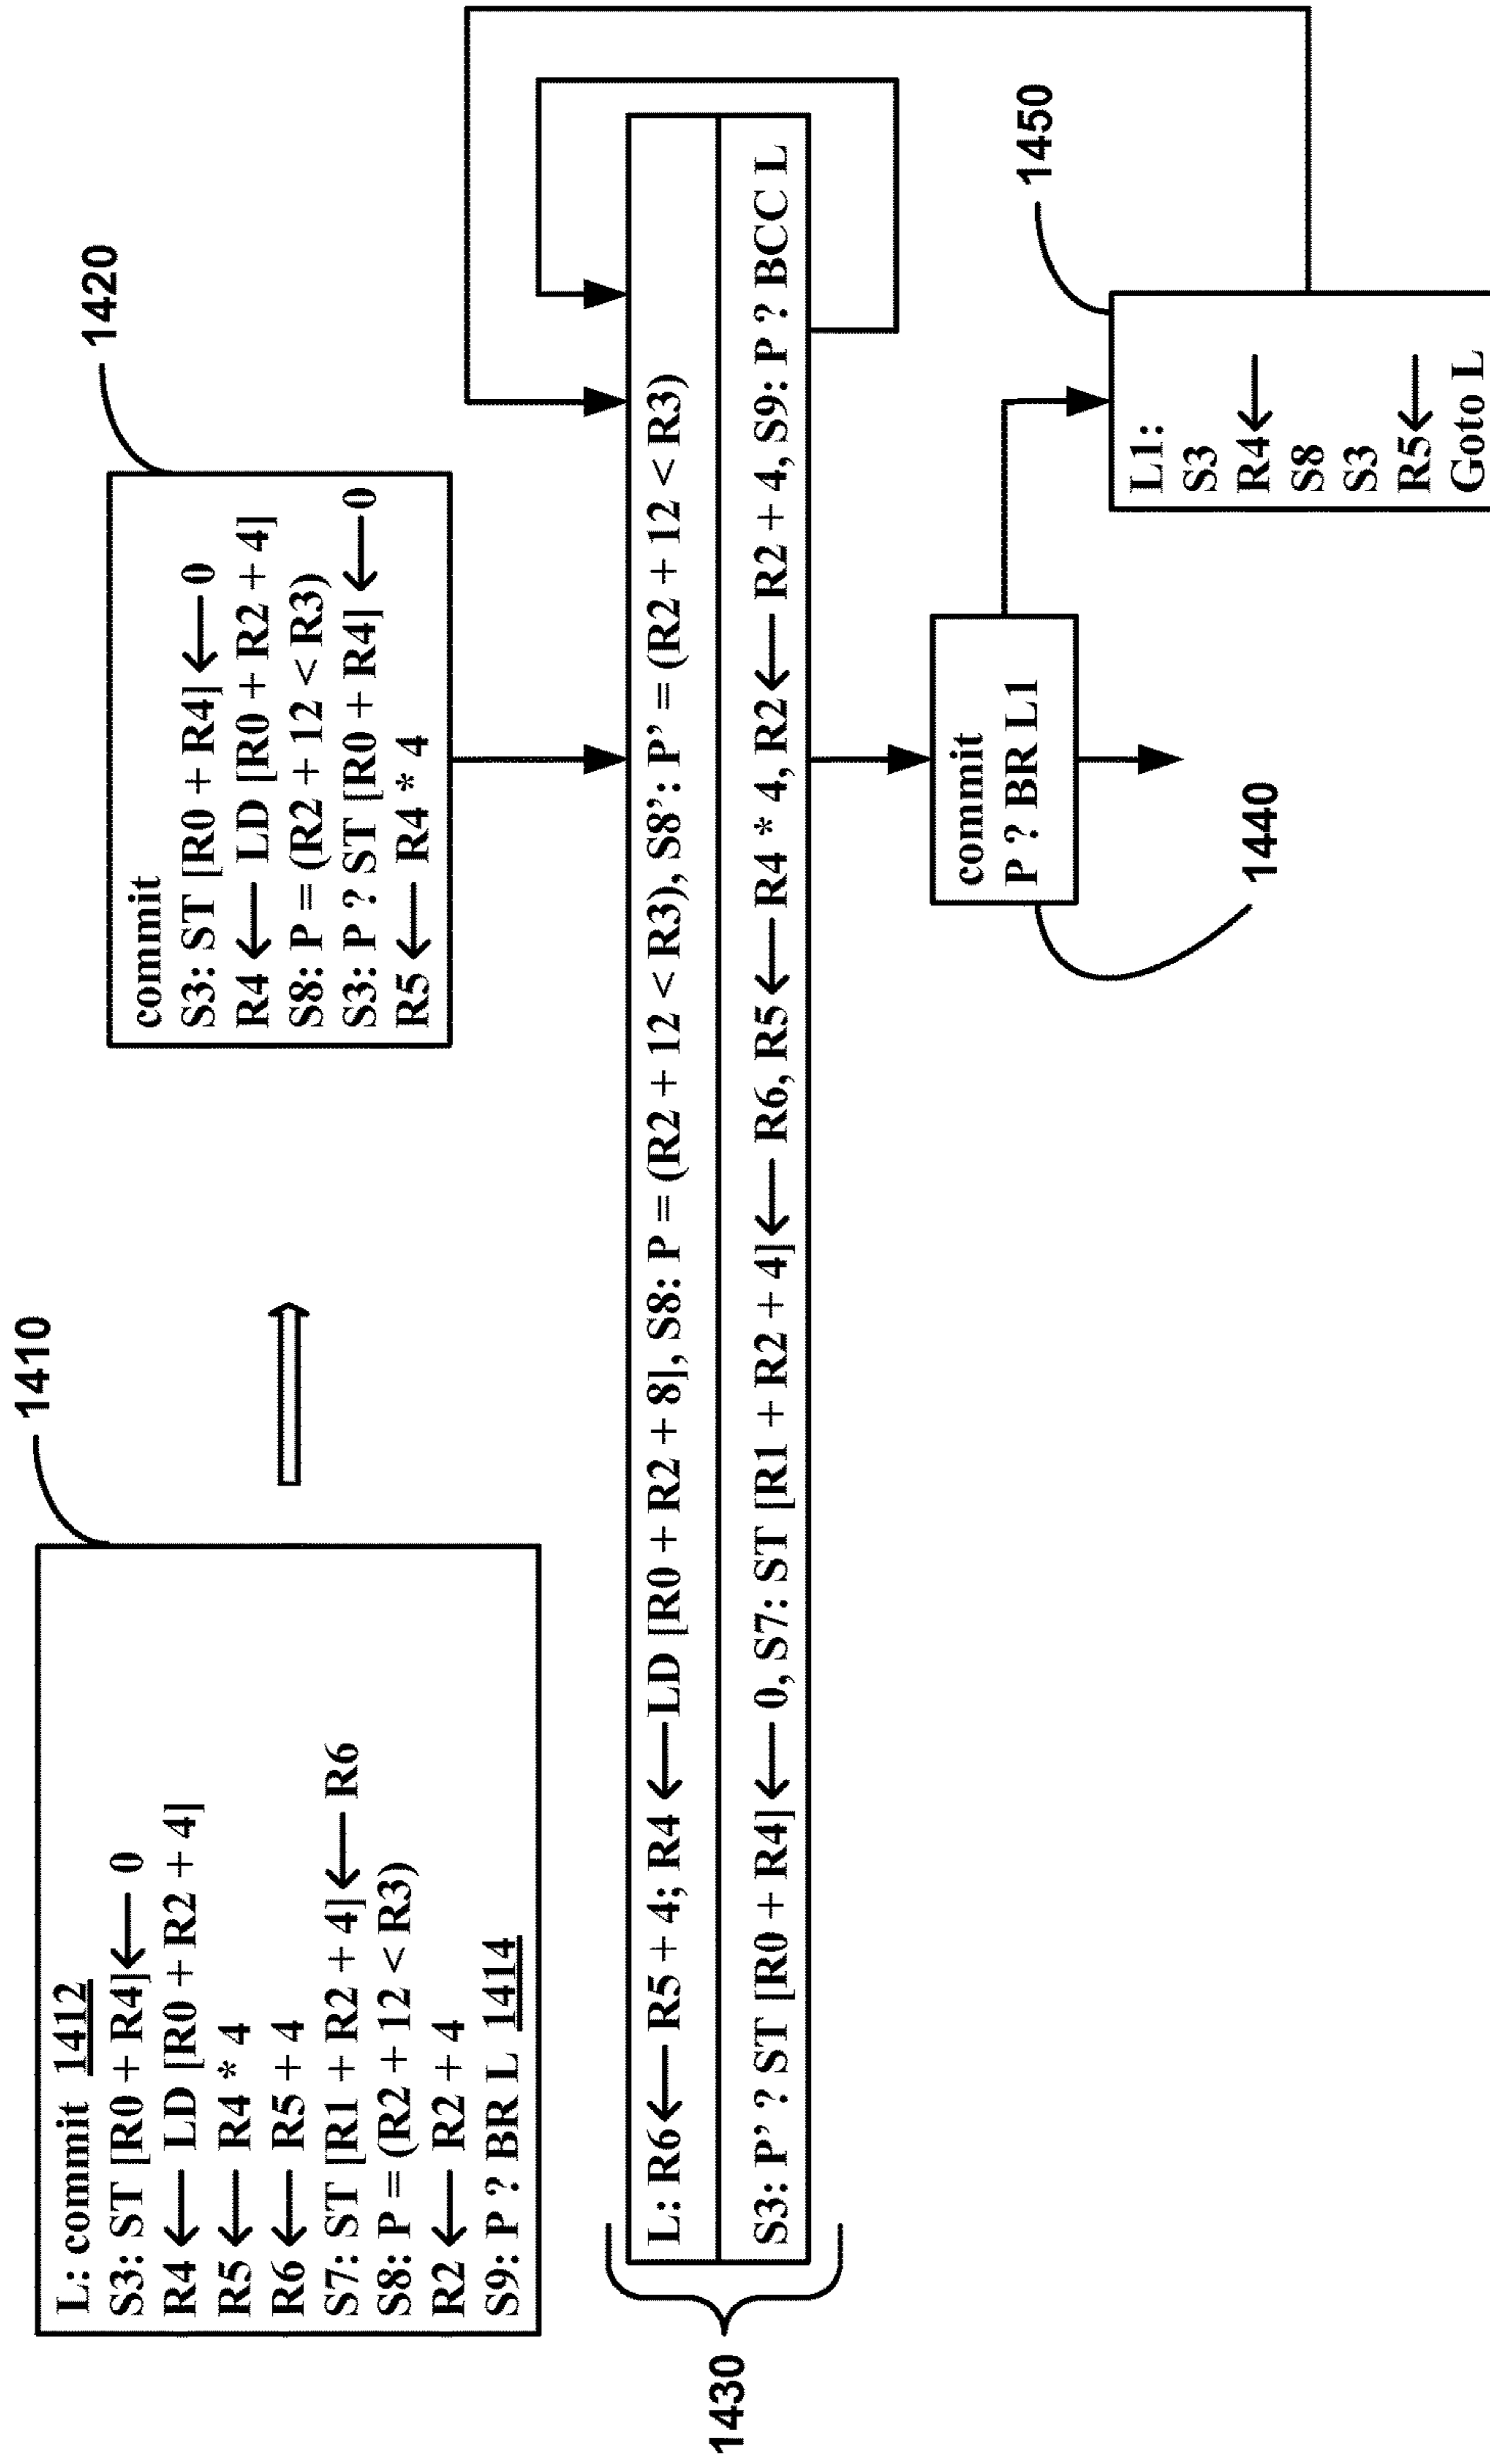


FIG. 14




FIG. 15

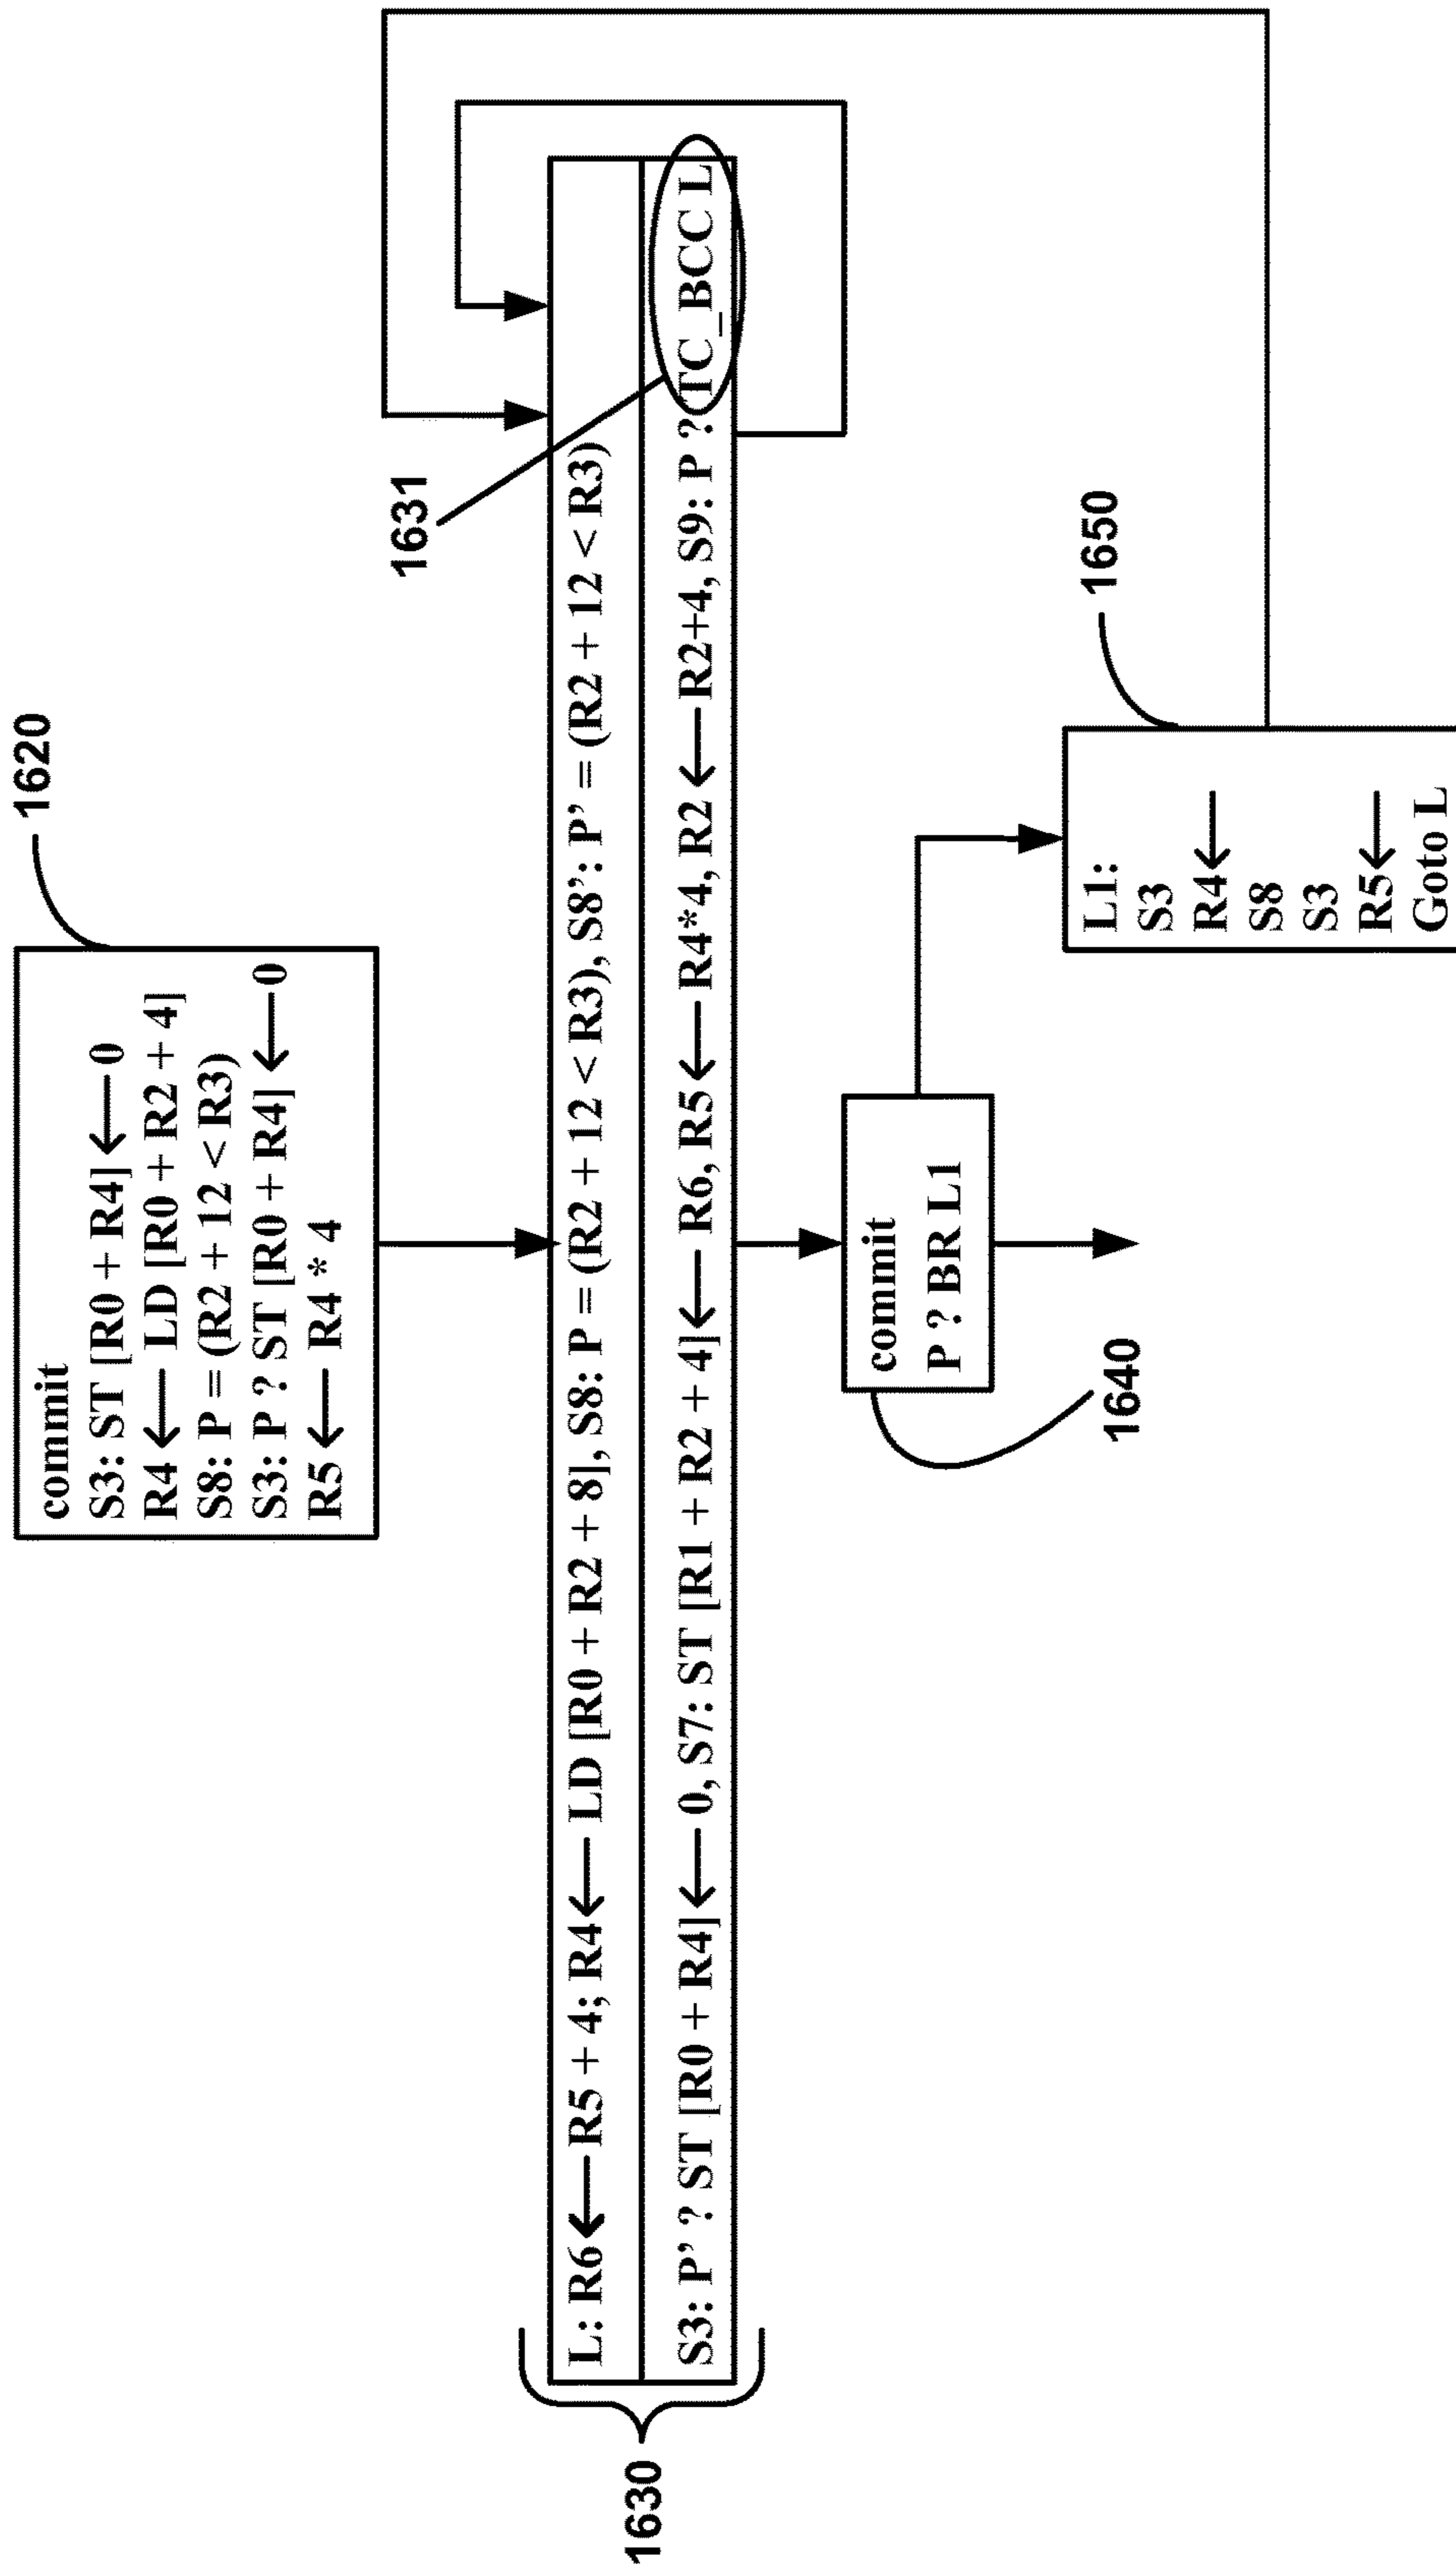


FIG. 16



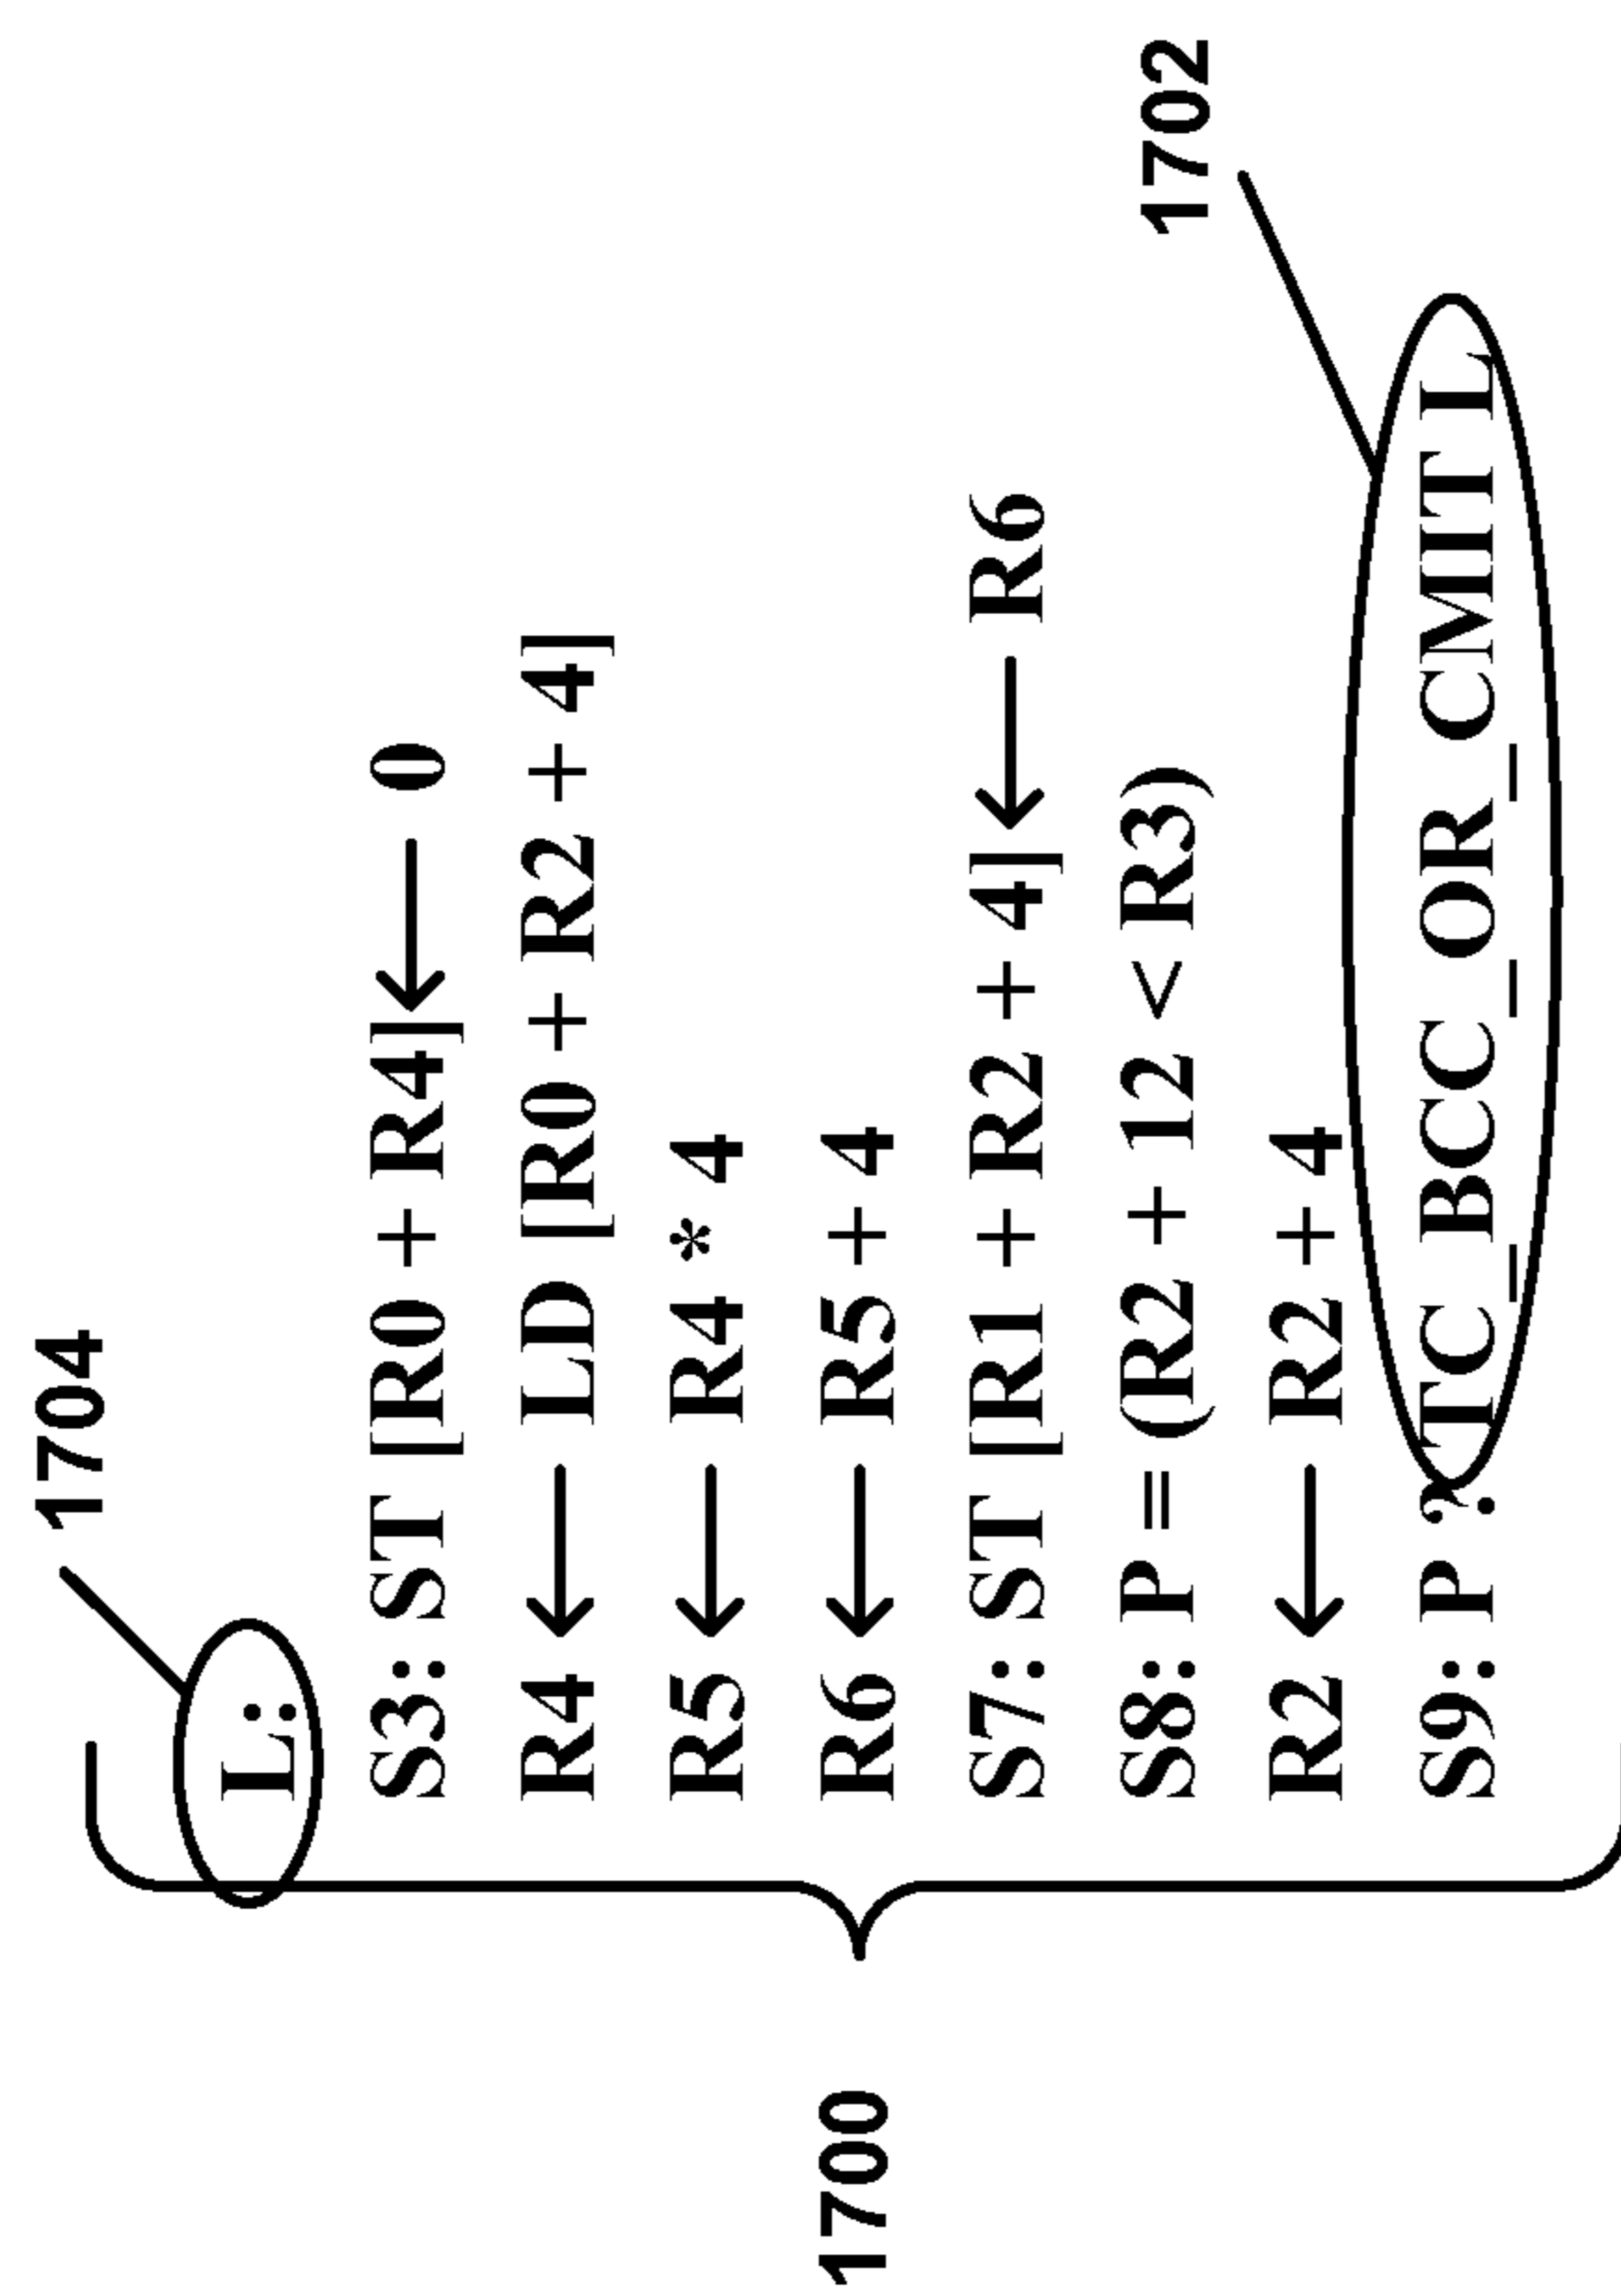


FIG. 17

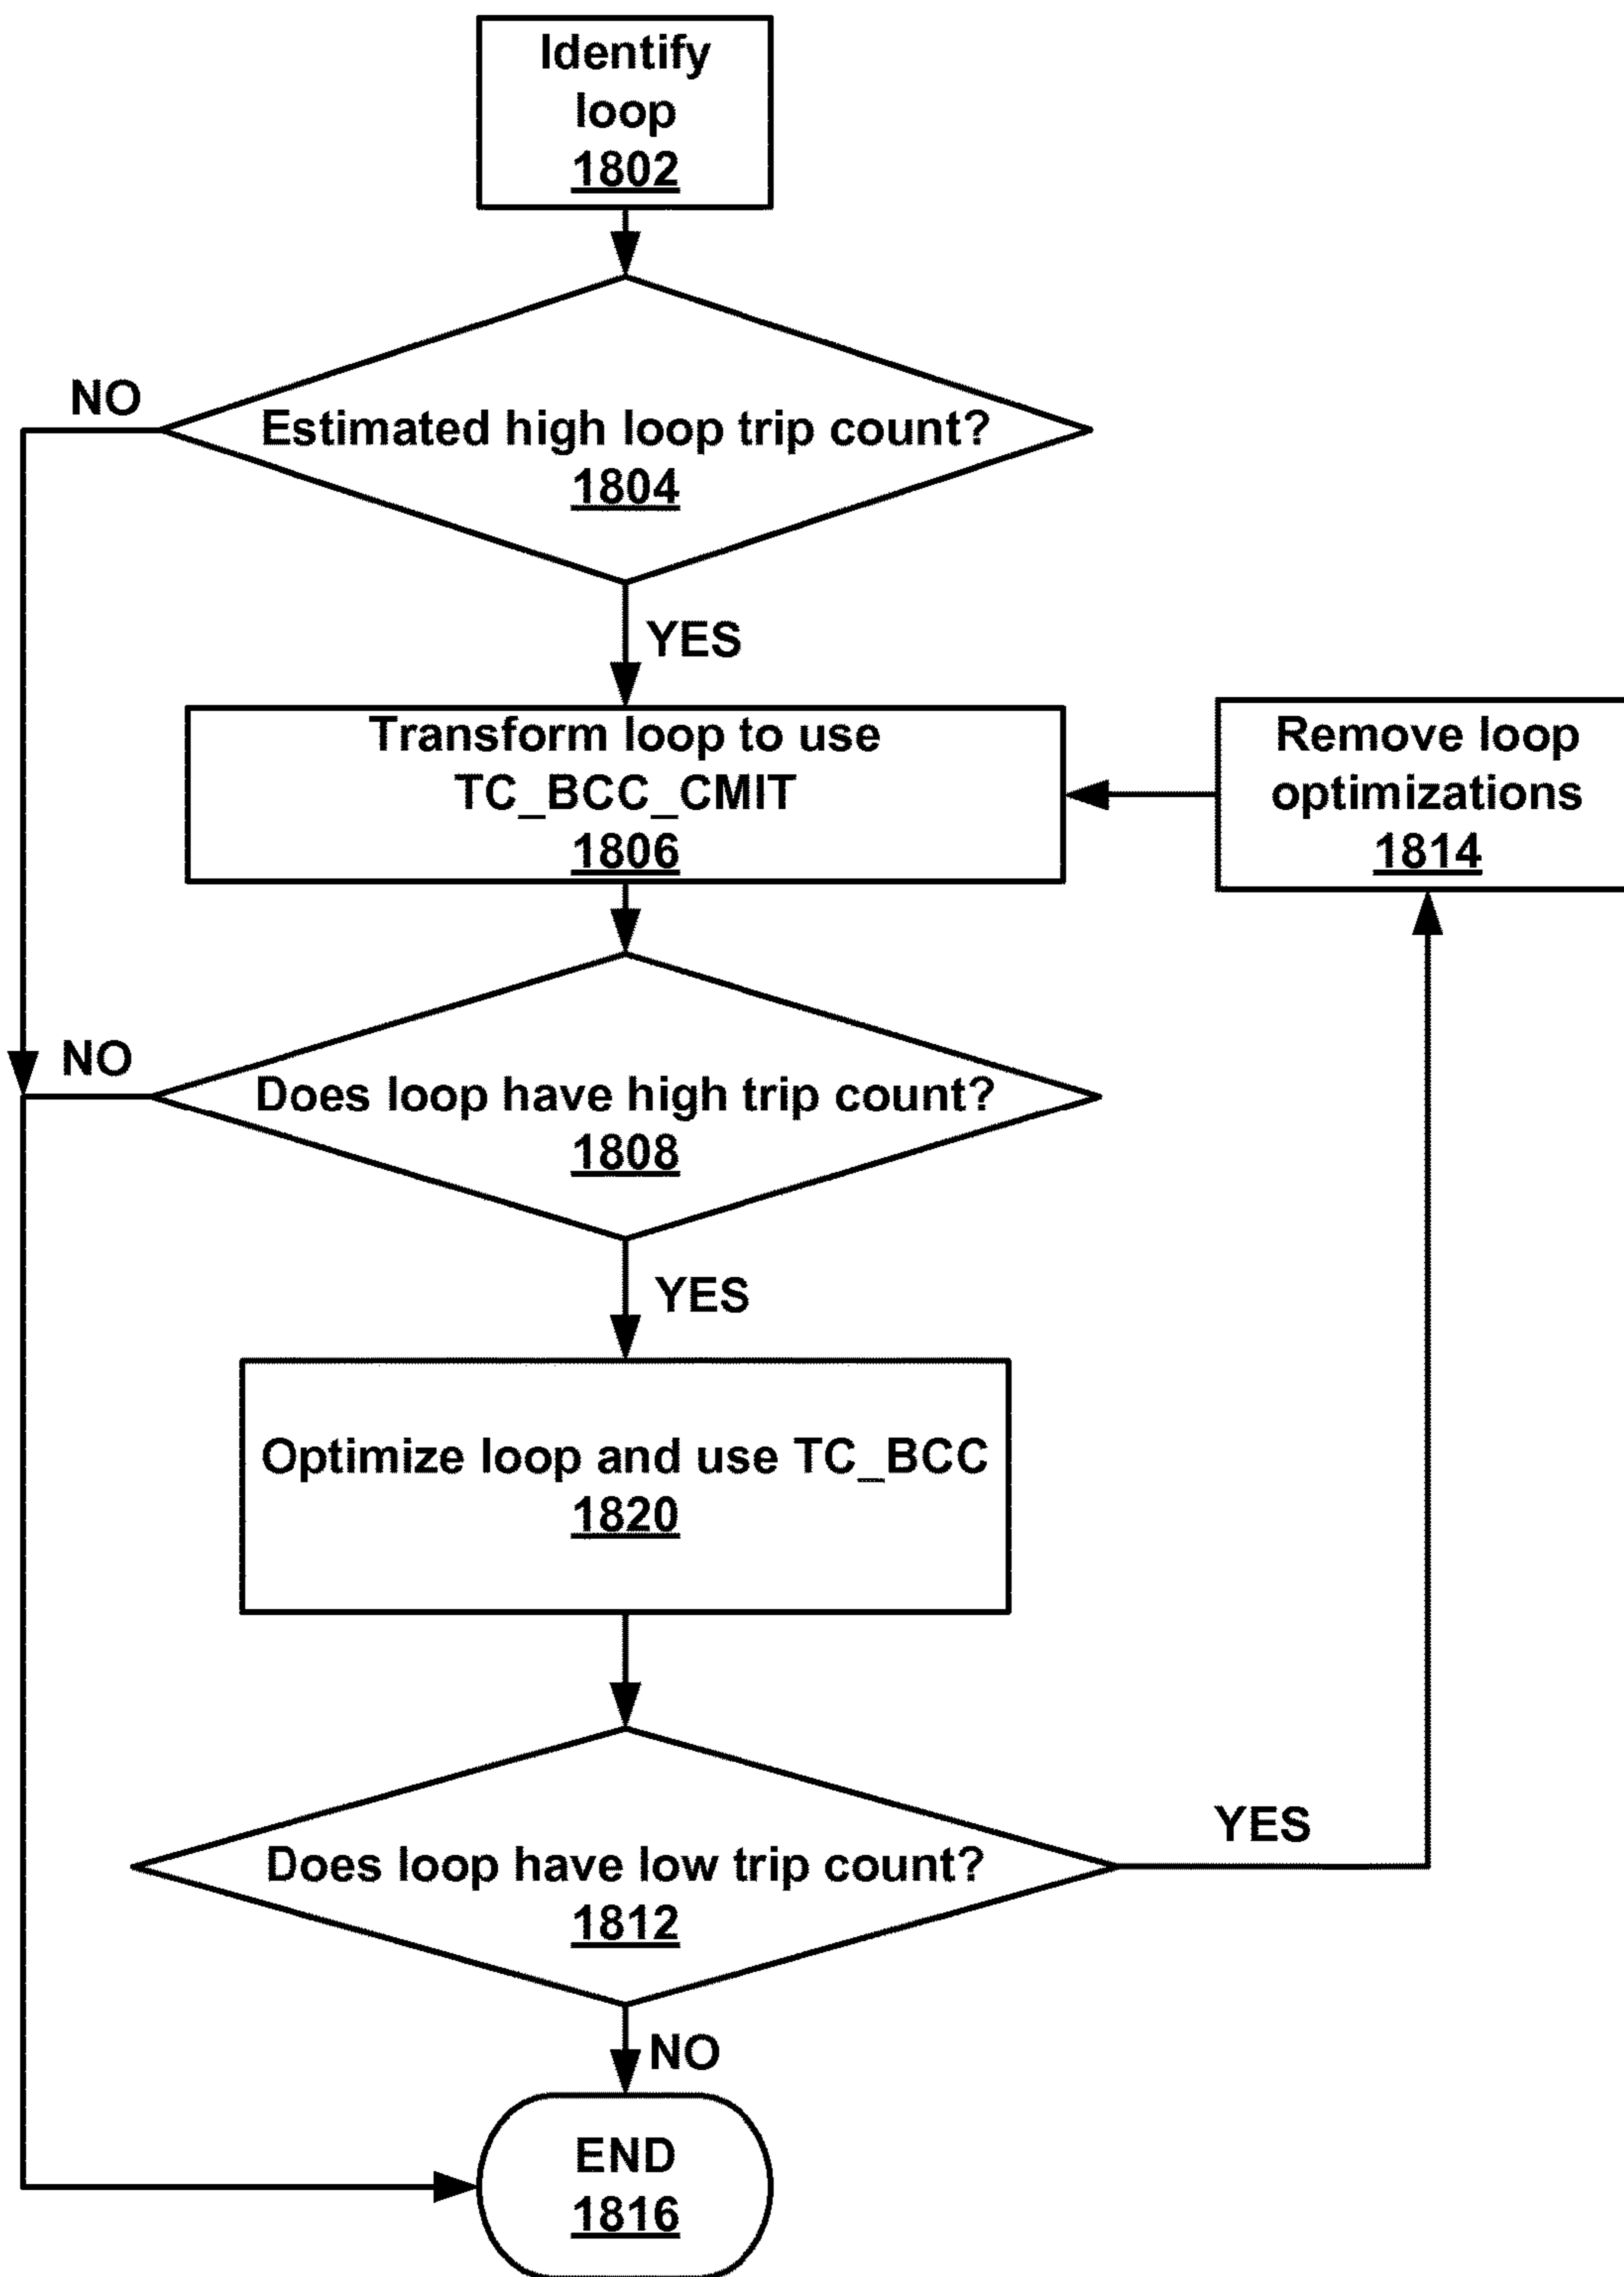


FIG. 18

## 1

# INSTRUCTION AND LOGIC TO MONITOR LOOP TRIP COUNT AND REMOVE LOOP OPTIMIZATIONS

## FIELD OF THE INVENTION

The present disclosure pertains to the field of processing logic, microprocessors, and associated instruction set architecture that, when executed by the processor or other processing logic, perform logical, mathematical, or other functional operations.

## DESCRIPTION OF RELATED ART

Utilizing computer coding algorithms, program code which is executed multiple times (such as a loop) can be optimized to increase performance. However, adding optimizations to loops may result in some fixed overhead which negatively impacts performance. Therefore, to benefit from optimization, the code within an optimized loop has to be executed a minimum number of times so that the performance benefit from optimization exceeds the negative impact on performance from the overhead. Traditionally, the number of times a loop is executed (also known as the loop trip count) was estimated by methods such as observing the execution time of a loop. Based on the estimate, loops were selectively optimized if the estimated loop trip count was high. However, such methods often do not provide an accurate estimate of the loop trip count, leading to unintended loop optimizations, and consequently, poor performance. Thus, there is a need to efficiently and accurately determine loop trip counts when program code is executed.

## DESCRIPTION OF THE FIGURES

Embodiments are illustrated by way of example and not limitation in the Figures of the accompanying drawings:

FIG. 1A is a block diagram of a system according to one embodiment;

FIG. 1B is a block diagram of a system according to one embodiment;

FIG. 1C is a block diagram of a system according to one embodiment;

FIG. 2 is a block diagram of a processor according to one embodiment;

FIG. 3A illustrates packed data types according to one embodiment;

FIG. 3B illustrates packed data types according one embodiment;

FIG. 3C illustrates packed data types according to one embodiment;

FIG. 3D illustrates an instruction encoding according to one embodiment;

FIG. 3E illustrates an instruction encoding according to one embodiment;

FIG. 3F illustrates an instruction encoding according to one embodiment;

FIG. 4A illustrates elements of a processor micro-architecture according to one embodiment;

FIG. 4B illustrates elements of a processor micro-architecture according to one embodiment;

FIG. 5 is a block diagram of a processor according to one embodiment;

FIG. 6 is a block diagram of a computer system according to one embodiment;

FIG. 7 is a block diagram of a computer system according to one embodiment;

## 2

FIG. 8 is a block diagram of a computer system according to one embodiment;

FIG. 9 is a block diagram of a system-on-a-chip according to one embodiment;

FIG. 10 is a block diagram of a processor according to one embodiment;

FIG. 11 is a block diagram of an IP core development system according to one embodiment;

FIG. 12 illustrates an architecture emulation system according to one embodiment.

FIG. 13 illustrates a system to translate instructions according to one embodiment.

FIG. 14 illustrates optimization of program code in an embodiment.

FIG. 15 illustrates exemplary optimized code to track loop trip count in an embodiment.

FIG. 16 illustrates exemplary optimized code to track loop trip count in an embodiment.

FIG. 17 illustrates exemplary code without loop optimizations which reduces commit overhead in an embodiment.

FIG. 18 illustrates a method for transforming code with loops in an embodiment.

## DETAILED DESCRIPTION

The following description describes an instruction and processing logic to efficiently monitor loop trip count within or in association with a processor, computer system, or other processing apparatus. In the following description, numerous specific details such as processing logic, processor types, micro-architectural conditions, events, enablement mechanisms, and the like are set forth in order to provide a more thorough understanding of embodiments of the present invention. It will be appreciated, however, by one skilled in the art that the invention may be practiced without such specific details. Additionally, some well known structures, circuits, and the like have not been shown in detail to avoid unnecessarily obscuring embodiments of the present invention.

Although the following embodiments are described with reference to a processor, other embodiments are applicable to other types of integrated circuits and logic devices. Similar techniques and teachings of embodiments of the present invention can be applied to other types of circuits or semiconductor devices that can benefit from higher pipeline throughput and improved performance. The teachings of embodiments of the present invention are applicable to any processor or machine that performs data manipulations. However, the present invention is not limited to processors or machines that perform 512 bit, 256 bit, 128 bit, 64 bit, 32 bit, or 16 bit data operations and can be applied to any processor and machine in which manipulation or management of data is performed. In addition, the following description provides examples, and the accompanying drawings show various examples for the purposes of illustration. However, these examples should not be construed in a limiting sense as they are merely intended to provide examples of embodiments of the present invention rather than to provide an exhaustive list of all possible implementations of embodiments of the present invention.

Although the below examples describe instruction handling and distribution in the context of execution units and logic circuits, other embodiments of the present invention can be accomplished by way of a data or instructions stored on a machine-readable, tangible medium, which when performed by a machine cause the machine to perform functions consistent with at least one embodiment of the inven-

tion. In one embodiment, functions associated with embodiments of the present invention are embodied in machine-executable instructions. The instructions can be used to cause a general-purpose or special-purpose processor that is programmed with the instructions to perform the steps of the present invention. Embodiments of the present invention may be provided as a computer program product or software which may include a machine or computer-readable medium having stored thereon instructions which may be used to program a computer (or other electronic devices) to perform one or more operations according to embodiments of the present invention. Alternatively, steps of embodiments of the present invention might be performed by specific hardware components that contain fixed-function logic for performing the steps, or by any combination of programmed computer components and fixed-function hardware components.

Instructions used to program logic to perform embodiments of the invention can be stored within a memory in the system, such as DRAM, cache, flash memory, or other storage. Furthermore, the instructions can be distributed via a network or by way of other computer readable media. Thus a machine-readable medium may include any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer), but is not limited to, floppy diskettes, optical disks, Compact Disc, Read-Only Memory (CD-ROMs), and magneto-optical disks, Read-Only Memory (ROMs), Random Access Memory (RAM), Erasable Programmable Read-Only Memory (EPROM), Electrically Erasable Programmable Read-Only Memory (EEPROM), magnetic or optical cards, flash memory, or a tangible, machine-readable storage used in the transmission of information over the Internet via electrical, optical, acoustical or other forms of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.). Accordingly, the computer-readable medium includes any type of tangible machine-readable medium suitable for storing or transmitting electronic instructions or information in a form readable by a machine (e.g., a computer).

A design may go through various stages, from creation to simulation to fabrication. Data representing a design may represent the design in a number of manners. First, as is useful in simulations, the hardware may be represented using a hardware description language or another functional description language. Additionally, a circuit level model with logic and/or transistor gates may be produced at some stages of the design process. Furthermore, most designs, at some stage, reach a level of data representing the physical placement of various devices in the hardware model. In the case where conventional semiconductor fabrication techniques are used, the data representing the hardware model may be the data specifying the presence or absence of various features on different mask layers for masks used to produce the integrated circuit. In any representation of the design, the data may be stored in any form of a machine readable medium. A memory or a magnetic or optical storage such as a disc may be the machine readable medium to store information transmitted via optical or electrical wave modulated or otherwise generated to transmit such information. When an electrical carrier wave indicating or carrying the code or design is transmitted, to the extent that copying, buffering, or re-transmission of the electrical signal is performed, a new copy is made. Thus, a communication provider or a network provider may store on a tangible, machine-readable medium, at least temporarily, an article, such as information encoded into a carrier wave, embodying techniques of embodiments of the present invention.

In modern processors, a number of different execution units are used to process and execute a variety of code and instructions. Not all instructions are created equal as some are quicker to complete while others can take a number of clock cycles to complete. The faster the throughput of instructions, the better the overall performance of the processor. Thus it would be advantageous to have as many instructions execute as fast as possible. However, there are certain instructions that have greater complexity and require more in terms of execution time and processor resources. For example, there are floating point instructions, load/store operations, data moves, etc.

As more computer systems are used in internet, text, and multimedia applications, additional processor support has been introduced over time. In one embodiment, an instruction set may be associated with one or more computer architectures, including data types, instructions, register architecture, addressing modes, memory architecture, interrupt and exception handling, and external input and output (I/O).

In one embodiment, the instruction set architecture (ISA) may be implemented by one or more micro-architectures, which includes processor logic and circuits used to implement one or more instruction sets. Accordingly, processors with different micro-architectures can share at least a portion of a common instruction set. For example, Intel® Pentium 4 processors, Intel® Core™ processors, and processors from Advanced Micro Devices, Inc. of Sunnyvale Calif. implement nearly identical versions of the x86 instruction set (with some extensions that have been added with newer versions), but have different internal designs. Similarly, processors designed by other processor development companies, such as ARM Holdings, Ltd., MIPS, or their licensees or adopters, may share at least a portion a common instruction set, but may include different processor designs. For example, the same register architecture of the ISA may be implemented in different ways in different micro-architectures using new or well-known techniques, including dedicated physical registers, one or more dynamically allocated physical registers using a register renaming mechanism (e.g., the use of a Register Alias Table (RAT), a Reorder Buffer (ROB) and a retirement register file. In one embodiment, registers may include one or more registers, register architectures, register files, or other register sets that may or may not be addressable by a software programmer.

In one embodiment, an instruction may include one or more instruction formats. In one embodiment, an instruction format may indicate various fields (number of bits, location of bits, etc.) to specify, among other things, the operation to be performed and the operand(s) on which that operation is to be performed. Some instruction formats may be further broken defined by instruction templates (or sub formats). For example, the instruction templates of a given instruction format may be defined to have different subsets of the instruction format's fields and/or defined to have a given field interpreted differently. In one embodiment, an instruction is expressed using an instruction format (and, if defined, in a given one of the instruction templates of that instruction format) and specifies or indicates the operation and the operands upon which the operation will operate.

Scientific, financial, auto-vectorized general purpose, RMS (recognition, mining, and synthesis), and visual and multimedia applications (e.g., 2D/3D graphics, image processing, video compression/decompression, voice recognition algorithms and audio manipulation) may require the same operation to be performed on a large number of data items. In one embodiment, Single Instruction Multiple Data

(SIMD) refers to a type of instruction that causes a processor to perform an operation on multiple data elements. SIMD technology may be used in processors that can logically divide the bits in a register into a number of fixed-sized or variable-sized data elements, each of which represents a separate value. For example, in one embodiment, the bits in a 64-bit register may be organized as a source operand containing four separate 16-bit data elements, each of which represents a separate 16-bit value. This type of data may be referred to as ‘packed’ data type or ‘vector’ data type, and operands of this data type are referred to as packed data operands or vector operands. In one embodiment, a packed data item or vector may be a sequence of packed data elements stored within a single register, and a packed data operand or a vector operand may be a source or destination operand of a SIMD instruction (or ‘packed data instruction’ or a ‘vector instruction’). In one embodiment, a SIMD instruction specifies a single vector operation to be performed on two source vector operands to generate a destination vector operand (also referred to as a result vector operand) of the same or different size, with the same or different number of data elements, and in the same or different data element order.

SIMD technology, such as that employed by the Intel® Core™ processors having an instruction set including x86, MMX™, Streaming SIMD Extensions (SSE), SSE2, SSE3, SSE4.1, and SSE4.2 instructions, ARM processors, such as the ARM Cortex® family of processors having an instruction set including the Vector Floating Point (VFP) and/or NEON instructions, and MIPS processors, such as the Loongson family of processors developed by the Institute of Computing Technology (ICT) of the Chinese Academy of Sciences, has enabled a significant improvement in application performance (Core™ and MMX™ are registered trademarks or trademarks of Intel Corporation of Santa Clara, Calif.).

In one embodiment, destination and source registers/data are generic terms to represent the source and destination of the corresponding data or operation. In some embodiments, they may be implemented by registers, memory, or other storage areas having other names or functions than those depicted. For example, in one embodiment, “DEST1” may be a temporary storage register or other storage area, whereas “SRC1” and “SRC2” may be a first and second source storage register or other storage area, and so forth. In other embodiments, two or more of the SRC and DEST storage areas may correspond to different data storage elements within the same storage area (e.g., a SIMD register). In one embodiment, one of the source registers may also act as a destination register by, for example, writing back the result of an operation performed on the first and second source data to one of the two source registers serving as a destination registers.

FIG. 1A is a block diagram of an exemplary computer system formed with a processor that includes execution units to execute an instruction in accordance with one embodiment of the present invention. System 100 includes a component, such as a processor 102 to employ execution units including logic to perform algorithms for process data, in accordance with the present invention, such as in the embodiment described herein. System 100 is representative of processing systems based on the PENTIUM® III, PENTIUM® 4, Xeon™, Itanium®, XScale™ and/or StrongARM™ microprocessors available from Intel Corporation of Santa Clara, Calif., although other systems (including PCs having other microprocessors, engineering workstations, set-top boxes and the like) may also be used. In one

embodiment, sample system 100 may execute a version of the WINDOWS™ operating system available from Microsoft Corporation of Redmond, Wash., although other operating systems (UNIX and Linux for example), embedded software, and/or graphical user interfaces, may also be used. Thus, embodiments of the present invention are not limited to any specific combination of hardware circuitry and software.

Embodiments are not limited to computer systems. Alternative embodiments of the present invention can be used in other devices such as handheld devices and embedded applications. Some examples of handheld devices include cellular phones, Internet Protocol devices, digital cameras, personal digital assistants (PDAs), and handheld PCs. Embedded applications can include a micro controller, a digital signal processor (DSP), system on a chip, network computers (NetPC), set-top boxes, network hubs, wide area network (WAN) switches, or any other system that can perform one or more instructions in accordance with at least one embodiment.

FIG. 1A is a block diagram of a computer system 100 formed with a processor 102 that includes one or more execution units 108 to perform an algorithm to perform at least one instruction in accordance with one embodiment of the present invention. One embodiment may be described in the context of a single processor desktop or server system, but alternative embodiments can be included in a multiprocessor system. System 100 is an example of a ‘hub’ system architecture. The computer system 100 includes a processor 102 to process data signals. The processor 102 can be a complex instruction set computer (CISC) microprocessor, a reduced instruction set computing (RISC) microprocessor, a very long instruction word (VLIW) microprocessor, a processor implementing a combination of instruction sets, or any other processor device, such as a digital signal processor, for example. The processor 102 is coupled to a processor bus 110 that can transmit data signals between the processor 102 and other components in the system 100. The elements of system 100 perform their conventional functions that are well known to those familiar with the art.

In one embodiment, the processor 102 includes a Level 1 (L1) internal cache memory 104. Depending on the architecture, the processor 102 can have a single internal cache or multiple levels of internal cache. Alternatively, in another embodiment, the cache memory can reside external to the processor 102. Other embodiments can also include a combination of both internal and external caches depending on the particular implementation and needs. Register file 106 can store different types of data in various registers including integer registers, floating point registers, status registers, and instruction pointer register.

Execution unit 108, including logic to perform integer and floating point operations, also resides in the processor 102. The processor 102 also includes a microcode (ucode) ROM that stores microcode for certain macroinstructions. For one embodiment, execution unit 108 includes logic to handle a packed instruction set 109. By including the packed instruction set 109 in the instruction set of a general-purpose processor 102, along with associated circuitry to execute the instructions, the operations used by many multimedia applications may be performed using packed data in a general-purpose processor 102. Thus, many multimedia applications can be accelerated and executed more efficiently by using the full width of a processor’s data bus for performing operations on packed data. This can eliminate the need to

transfer smaller units of data across the processor's data bus to perform one or more operations one data element at a time.

Alternate embodiments of an execution unit **108** can also be used in micro controllers, embedded processors, graphics devices, DSPs, and other types of logic circuits. System **100** includes a memory **120**. Memory **120** can be a dynamic random access memory (DRAM) device, a static random access memory (SRAM) device, flash memory device, or other memory device. Memory **120** can store instructions and/or data represented by data signals that can be executed by the processor **102**.

A system logic chip **116** is coupled to the processor bus **110** and memory **120**. The system logic chip **116** in the illustrated embodiment is a memory controller hub (MCH). The processor **102** can communicate to the MCH **116** via a processor bus **110**. The MCH **116** provides a high bandwidth memory path **118** to memory **120** for instruction and data storage and for storage of graphics commands, data and textures. The MCH **116** is to direct data signals between the processor **102**, memory **120**, and other components in the system **100** and to bridge the data signals between processor bus **110**, memory **120**, and system I/O **122**. In some embodiments, the system logic chip **116** can provide a graphics port for coupling to a graphics controller **112**. The MCH **116** is coupled to memory **120** through a memory interface **118**. The graphics card **112** is coupled to the MCH **116** through an Accelerated Graphics Port (AGP) interconnect **114**.

System **100** uses a proprietary hub interface bus **122** to couple the MCH **116** to the I/O controller hub (ICH) **130**. The ICH **130** provides direct connections to some I/O devices via a local I/O bus. The local I/O bus is a high-speed I/O bus for connecting peripherals to the memory **120**, chipset, and processor **102**. Some examples are the audio controller, firmware hub (flash BIOS) **128**, wireless transceiver **126**, data storage **124**, legacy I/O controller containing user input and keyboard interfaces, a serial expansion port such as Universal Serial Bus (USB), and a network controller **134**. The data storage device **124** can comprise a hard disk drive, a floppy disk drive, a CD-ROM device, a flash memory device, or other mass storage device.

For another embodiment of a system, an instruction in accordance with one embodiment can be used with a system on a chip. One embodiment of a system on a chip comprises of a processor and a memory. The memory for one such system is a flash memory. The flash memory can be located on the same die as the processor and other system components. Additionally, other logic blocks such as a memory controller or graphics controller can also be located on a system on a chip.

FIG. 1B illustrates a data processing system **140** which implements the principles of one embodiment of the present invention. It will be readily appreciated by one of skill in the art that the embodiments described herein can be used with alternative processing systems without departure from the scope of embodiments of the invention.

Computer system **140** comprises a processing core **159** capable of performing at least one instruction in accordance with one embodiment. For one embodiment, processing core **159** represents a processing unit of any type of architecture, including but not limited to a CISC, a RISC or a VLIW type architecture. Processing core **159** may also be suitable for manufacture in one or more process technologies and by being represented on a machine readable media in sufficient detail, may be suitable to facilitate said manufacture.

Processing core **159** comprises an execution unit **142**, a set of register file(s) **145**, and a decoder **144**. Processing core

**159** also includes additional circuitry (not shown) which is not necessary to the understanding of embodiments of the present invention. Execution unit **142** is used for executing instructions received by processing core **159**. In addition to performing typical processor instructions, execution unit **142** can perform instructions in packed instruction set **143** for performing operations on packed data formats. Packed instruction set **143** includes instructions for performing embodiments of the invention and other packed instructions. Execution unit **142** is coupled to register file **145** by an internal bus. Register file **145** represents a storage area on processing core **159** for storing information, including data. As previously mentioned, it is understood that the storage area used for storing the packed data is not critical. Execution unit **142** is coupled to decoder **144**. Decoder **144** is used for decoding instructions received by processing core **159** into control signals and/or microcode entry points. In response to these control signals and/or microcode entry points, execution unit **142** performs the appropriate operations. In one embodiment, the decoder is used to interpret the opcode of the instruction, which will indicate what operation should be performed on the corresponding data indicated within the instruction.

Processing core **159** is coupled with bus **141** for communicating with various other system devices, which may include but are not limited to, for example, synchronous dynamic random access memory (SDRAM) control **146**, static random access memory (SRAM) control **147**, burst flash memory interface **148**, personal computer memory card international association (PCMCIA)/compact flash (CF) card control **149**, liquid crystal display (LCD) control **150**, direct memory access (DMA) controller **151**, and alternative bus master interface **152**. In one embodiment, data processing system **140** may also comprise an I/O bridge **154** for communicating with various I/O devices via an I/O bus **153**. Such I/O devices may include but are not limited to, for example, universal asynchronous receiver/transmitter (UART) **155**, universal serial bus (USB) **156**, Bluetooth wireless UART **157** and I/O expansion interface **158**.

One embodiment of data processing system **140** provides for mobile, network and/or wireless communications and a processing core **159** capable of performing SIMD operations including a text string comparison operation. Processing core **159** may be programmed with various audio, video, imaging and communications algorithms including discrete transformations such as a Walsh-Hadamard transform, a fast Fourier transform (FFT), a discrete cosine transform (DCT), and their respective inverse transforms; compression/decompression techniques such as color space transformation, video encode motion estimation or video decode motion compensation; and modulation/demodulation (MODEM) functions such as pulse coded modulation (PCM).

FIG. 1C illustrates yet alternative embodiments of a data processing system capable of performing SIMD text string comparison operations. In accordance with one alternative embodiment, data processing system **160** may include a main processor **166**, a SIMD coprocessor **161**, a cache memory **167**, and an input/output system **168**. The input/output system **168** may optionally be coupled to a wireless interface **169**. SIMD coprocessor **161** is capable of performing operations including instructions in accordance with one embodiment. Processing core **170** may be suitable for manufacture in one or more process technologies and by being represented on a machine readable media in sufficient detail, may be suitable to facilitate the manufacture of all or part of data processing system **160** including processing core **170**.

For one embodiment, SIMD coprocessor **161** comprises an execution unit **162** and a set of register file(s) **164**. One embodiment of main processor **165** comprises a decoder **165** to recognize instructions of instruction set **163** including instructions in accordance with one embodiment for execution by execution unit **162**. For alternative embodiments, SIMD coprocessor **161** also comprises at least part of decoder **165B** to decode instructions of instruction set **163**. Processing core **170** also includes additional circuitry (not shown) which is not necessary to the understanding of embodiments of the present invention.

In operation, the main processor **166** executes a stream of data processing instructions that control data processing operations of a general type including interactions with the cache memory **167**, and the input/output system **168**. Embedded within the stream of data processing instructions are SIMD coprocessor instructions. The decoder **165** of main processor **166** recognizes these SIMD coprocessor instructions as being of a type that should be executed by an attached SIMD coprocessor **161**. Accordingly, the main processor **166** issues these SIMD coprocessor instructions (or control signals representing SIMD coprocessor instructions) on the coprocessor bus **172** where from they are received by any attached SIMD coprocessors. In this case, the SIMD coprocessor **161** will accept and execute any received SIMD coprocessor instructions intended for it.

Data may be received via wireless interface **169** for processing by the SIMD coprocessor instructions. For one example, voice communication may be received in the form of a digital signal, which may be processed by the SIMD coprocessor instructions to regenerate digital audio samples representative of the voice communications. For another example, compressed audio and/or video may be received in the form of a digital bit stream, which may be processed by the SIMD coprocessor instructions to regenerate digital audio samples and/or motion video frames. For one embodiment of processing core **170**, main processor **166**, and a SIMD coprocessor **161** are integrated into a single processing core **170** comprising an execution unit **162**, a set of register file(s) **164**, and a decoder **165** to recognize instructions of instruction set **163** including instructions in accordance with one embodiment.

FIG. **2** is a block diagram of the micro-architecture for a processor **200** that includes logic circuits to perform instructions in accordance with one embodiment of the present invention. In some embodiments, an instruction in accordance with one embodiment can be implemented to operate on data elements having sizes of byte, word, doubleword, quadword, etc., as well as datatypes, such as single and double precision integer and floating point datatypes. In one embodiment the in-order front end **201** is the part of the processor **200** that fetches instructions to be executed and prepares them to be used later in the processor pipeline. The front end **201** may include several units. In one embodiment, the instruction prefetcher **226** fetches instructions from memory and feeds them to an instruction decoder **228** which in turn decodes or interprets them. For example, in one embodiment, the decoder decodes a received instruction into one or more operations called "micro-instructions" or "micro-operations" (also called micro op or uops) that the machine can execute. In other embodiments, the decoder parses the instruction into an opcode and corresponding data and control fields that are used by the micro-architecture to perform operations in accordance with one embodiment. In one embodiment, the trace cache **230** takes decoded uops and assembles them into program ordered sequences or traces in the uop queue **234** for execution. When the trace

cache **230** encounters a complex instruction, the microcode ROM **232** provides the uops needed to complete the operation.

Some instructions are converted into a single micro-op, whereas others need several micro-ops to complete the full operation. In one embodiment, if more than four micro-ops are needed to complete a instruction, the decoder **228** accesses the microcode ROM **232** to do the instruction. For one embodiment, an instruction can be decoded into a small number of micro ops for processing at the instruction decoder **228**. In another embodiment, an instruction can be stored within the microcode ROM **232** should a number of micro-ops be needed to accomplish the operation. The trace cache **230** refers to a entry point programmable logic array (PLA) to determine a correct micro-instruction pointer for reading the micro-code sequences to complete one or more instructions in accordance with one embodiment from the micro-code ROM **232**. After the microcode ROM **232** finishes sequencing micro-ops for an instruction, the front end **201** of the machine resumes fetching micro-ops from the trace cache **230**.

The out-of-order execution engine **203** is where the instructions are prepared for execution. The out-of-order execution logic has a number of buffers to smooth out and re-order the flow of instructions to optimize performance as they go down the pipeline and get scheduled for execution. The allocator logic allocates the machine buffers and resources that each uop needs in order to execute. The register renaming logic renames logic registers onto entries in a register file. The allocator also allocates an entry for each uop in one of the two uop queues, one for memory operations and one for non-memory operations, in front of the instruction schedulers: memory scheduler, fast scheduler **202**, slow/general floating point scheduler **204**, and simple floating point scheduler **206**. The uop schedulers **202**, **204**, **206**, determine when a uop is ready to execute based on the readiness of their dependent input register operand sources and the availability of the execution resources the uops need to complete their operation. The fast scheduler **202** of one embodiment can schedule on each half of the main clock cycle while the other schedulers can only schedule once per main processor clock cycle. The schedulers arbitrate for the dispatch ports to schedule uops for execution.

Register files **208**, **210**, sit between the schedulers **202**, **204**, **206**, and the execution units **212**, **214**, **216**, **218**, **220**, **222**, **224** in the execution block **211**. There is a separate register file **208**, **210**, for integer and floating point operations, respectively. Each register file **208**, **210**, of one embodiment also includes a bypass network that can bypass or forward just completed results that have not yet been written into the register file to new dependent uops. The integer register file **208** and the floating point register file **210** are also capable of communicating data with the other. For one embodiment, the integer register file **208** is split into two separate register files, one register file for the low order 32 bits of data and a second register file for the high order 32 bits of data. The floating point register file **210** of one embodiment has 128 bit wide entries because floating point instructions typically have operands from 64 to 128 bits in width.

The execution block **211** contains the execution units **212**, **214**, **216**, **218**, **220**, **222**, **224**, where the instructions are actually executed. This section includes the register files **208**, **210**, that store the integer and floating point data operand values that the micro-instructions need to execute. The processor **200** of one embodiment is comprised of a number of execution units: address generation unit (AGU)

212, AGU 214, fast ALU 216, fast ALU 218, slow ALU 220, floating point ALU 222, floating point move unit 224. For one embodiment, the floating point execution blocks 222, 224, execute floating point, MMX, SIMD, and SSE, or other operations. The floating point ALU 222 of one embodiment includes a 64 bit by 64 bit floating point divider to execute divide, square root, and remainder micro-ops. For embodiments of the present invention, instructions involving a floating point value may be handled with the floating point hardware. In one embodiment, the ALU operations go to the high-speed ALU execution units 216, 218. The fast ALUs 216, 218, of one embodiment can execute fast operations with an effective latency of half a clock cycle. For one embodiment, most complex integer operations go to the slow ALU 220 as the slow ALU 220 includes integer execution hardware for long latency type of operations, such as a multiplier, shifts, flag logic, and branch processing. Memory load/store operations are executed by the AGUs 212, 214. For one embodiment, the integer ALUs 216, 218, 220, are described in the context of performing integer operations on 64 bit data operands. In alternative embodiments, the ALUs 216, 218, 220, can be implemented to support a variety of data bits including 16, 32, 128, 256, etc. Similarly, the floating point units 222, 224, can be implemented to support a range of operands having bits of various widths. For one embodiment, the floating point units 222, 224, can operate on 128 bits wide packed data operands in conjunction with SIMD and multimedia instructions.

In one embodiment, the uops schedulers 202, 204, 206, dispatch dependent operations before the parent load has finished executing. As uops are speculatively scheduled and executed in processor 200, the processor 200 also includes logic to handle memory misses. If a data load misses in the data cache, there can be dependent operations in flight in the pipeline that have left the scheduler with temporarily incorrect data. A replay mechanism tracks and re-executes instructions that use incorrect data. Only the dependent operations need to be replayed and the independent ones are allowed to complete. The schedulers and replay mechanism of one embodiment of a processor are also designed to catch instruction sequences for text string comparison operations.

The term “registers” may refer to the on-board processor storage locations that are used as part of instructions to identify operands. In other words, registers may be those that are usable from the outside of the processor (from a programmer’s perspective). However, the registers of an embodiment should not be limited in meaning to a particular type of circuit. Rather, a register of an embodiment is capable of storing and providing data, and performing the functions described herein. The registers described herein can be implemented by circuitry within a processor using any number of different techniques, such as dedicated physical registers, dynamically allocated physical registers using register renaming, combinations of dedicated and dynamically allocated physical registers, etc. In one embodiment, integer registers store thirty-two bit integer data. A register file of one embodiment also contains eight multimedia SIMD registers for packed data. For the discussions below, the registers are understood to be data registers designed to hold packed data, such as 64 bits wide MMX™ registers (also referred to as ‘mm’ registers in some instances) in microprocessors enabled with MMX technology from Intel Corporation of Santa Clara, Calif. These MMX registers, available in both integer and floating point forms, can operate with packed data elements that accompany SIMD and SSE instructions. Similarly, 128 bits wide XMM registers relating to SSE2, SSE3, SSE4, or beyond (referred to

generically as “SSEx”) technology can also be used to hold such packed data operands. In one embodiment, in storing packed data and integer data, the registers do not need to differentiate between the two data types. In one embodiment, integer and floating point are either contained in the same register file or different register files. Furthermore, in one embodiment, floating point and integer data may be stored in different registers or the same registers.

In the examples of the following figures, a number of data operands are described. FIG. 3A illustrates various packed data type representations in multimedia registers according to one embodiment of the present invention. FIG. 3A illustrates data types for a packed byte 310, a packed word 320, and a packed doubleword (dword) 330 for 128 bits wide operands. The packed byte format 310 of this example is 128 bits long and contains sixteen packed byte data elements. A byte is defined here as 8 bits of data. Information for each byte data element is stored in bit 7 through bit 0 for byte 0, bit 15 through bit 8 for byte 1, bit 23 through bit 16 for byte 2, and finally bit 120 through bit 127 for byte 15. Thus, all available bits are used in the register. This storage arrangement increases the storage efficiency of the processor. As well, with sixteen data elements accessed, one operation can now be performed on sixteen data elements in parallel.

Generally, a data element is an individual piece of data that is stored in a single register or memory location with other data elements of the same length. In packed data sequences relating to SSEx technology, the number of data elements stored in a XMM register is 128 bits divided by the length in bits of an individual data element. Similarly, in packed data sequences relating to MMX and SSE technology, the number of data elements stored in an MMX register is 64 bits divided by the length in bits of an individual data element. Although the data types illustrated in FIG. 3A are 128 bit long, embodiments of the present invention can also operate with 64 bit wide or other sized operands. The packed word format 320 of this example is 128 bits long and contains eight packed word data elements. Each packed word contains sixteen bits of information. The packed doubleword format 330 of FIG. 3A is 128 bits long and contains four packed doubleword data elements. Each packed doubleword data element contains thirty two bits of information. A packed quadword is 128 bits long and contains two packed quad-word data elements.

FIG. 3B illustrates alternative in-register data storage formats. Each packed data can include more than one independent data element. Three packed data formats are illustrated; packed half 341, packed single 342, and packed double 343. One embodiment of packed half 341, packed single 342, and packed double 343 contain fixed-point data elements. For an alternative embodiment one or more of packed half 341, packed single 342, and packed double 343 may contain floating-point data elements. One alternative embodiment of packed half 341 is one hundred twenty-eight bits long containing eight 16-bit data elements. One embodiment of packed single 342 is one hundred twenty-eight bits long and contains four 32-bit data elements. One embodiment of packed double 343 is one hundred twenty-eight bits long and contains two 64-bit data elements. It will be appreciated that such packed data formats may be further extended to other register lengths, for example, to 96-bits, 160-bits, 192-bits, 224-bits, 256-bits or more.

FIG. 3C illustrates various signed and unsigned packed data type representations in multimedia registers according to one embodiment of the present invention. Unsigned packed byte representation 344 illustrates the storage of an



unsigned packed byte in a SIMD register. Information for each byte data element is stored in bit seven through bit zero for byte zero, bit fifteen through bit eight for byte one, bit twenty-three through bit sixteen for byte two, and finally bit one hundred twenty through bit one hundred twenty-seven for byte fifteen. Thus, all available bits are used in the register. This storage arrangement can increase the storage efficiency of the processor. As well, with sixteen data elements accessed, one operation can now be performed on sixteen data elements in a parallel fashion. Signed packed byte representation **345** illustrates the storage of a signed packed byte. Note that the eighth bit of every byte data element is the sign indicator. Unsigned packed word representation **346** illustrates how word seven through word zero are stored in a SIMD register. Signed packed word representation **347** is similar to the unsigned packed word in-register representation **346**. Note that the sixteenth bit of each word data element is the sign indicator. Unsigned packed doubleword representation **348** shows how doubleword data elements are stored. Signed packed doubleword representation **349** is similar to unsigned packed doubleword in-register representation **348**. Note that the necessary sign bit is the thirty-second bit of each doubleword data element.

FIG. 3D is a depiction of one embodiment of an operation encoding (opcode) format **360**, having thirty-two or more bits, and register/memory operand addressing modes corresponding with a type of opcode format described in the "IA-32 Intel Architecture Software Developer's Manual Volume 2: Instruction Set Reference," which is available from Intel Corporation, Santa Clara, Calif. on the world-wide-web (www) at intel.com/design/litcentr. In one embodiment, an instruction may be encoded by one or more of fields **361** and **362**. Up to two operand locations per instruction may be identified, including up to two source operand identifiers **364** and **365**. For one embodiment, destination operand identifier **366** is the same as source operand identifier **364**, whereas in other embodiments they are different. For an alternative embodiment, destination operand identifier **366** is the same as source operand identifier **365**, whereas in other embodiments they are different. In one embodiment, one of the source operands identified by source operand identifiers **364** and **365** is overwritten by the results of the text string comparison operations, whereas in other embodiments identifier **364** corresponds to a source register element and identifier **365** corresponds to a destination register element. For one embodiment, operand identifiers **364** and **365** may be used to identify 32-bit or 64-bit source and destination operands.

FIG. 3E is a depiction of another alternative operation encoding (opcode) format **370**, having forty or more bits. Opcode format **370** corresponds with opcode format **360** and comprises an optional prefix byte **378**. An instruction according to one embodiment may be encoded by one or more of fields **378**, **371**, and **372**. Up to two operand locations per instruction may be identified by source operand identifiers **374** and **375** and by prefix byte **378**. For one embodiment, prefix byte **378** may be used to identify 32-bit or 64-bit source and destination operands. For one embodiment, destination operand identifier **376** is the same as source operand identifier **374**, whereas in other embodiments they are different. For an alternative embodiment, destination operand identifier **376** is the same as source operand identifier **375**, whereas in other embodiments they are different. In one embodiment, an instruction operates on one or more of the operands identified by operand identifiers **374** and **375** and one or more operands identified by the operand identifiers **374** and **375** is overwritten by the results

of the instruction, whereas in other embodiments, operands identified by identifiers **374** and **375** are written to another data element in another register. Opcode formats **360** and **370** allow register to register, memory to register, register by memory, register by register, register by immediate, register to memory addressing specified in part by MOD fields **363** and **373** and by optional scale-index-base and displacement bytes.

Turning next to FIG. 3F, in some alternative embodiments, 64 bit single instruction multiple data (SIMD) arithmetic operations may be performed through a coprocessor data processing (CDP) instruction. Operation encoding (opcode) format **380** depicts one such CDP instruction having CDP opcode fields **382** and **389**. The type of CDP instruction, for alternative embodiments, operations may be encoded by one or more of fields **383**, **384**, **387**, and **388**. Up to three operand locations per instruction may be identified, including up to two source operand identifiers **385** and **390** and one destination operand identifier **386**. One embodiment of the coprocessor can operate on 8, 16, 32, and 64 bit values. For one embodiment, an instruction is performed on integer data elements. In some embodiments, an instruction may be executed conditionally, using condition field **381**. For some embodiments, source data sizes may be encoded by field **383**. In some embodiments, Zero (Z), negative (N), carry (C), and overflow (V) detection can be done on SIMD fields. For some instructions, the type of saturation may be encoded by field **384**.

FIG. 4A is a block diagram illustrating an in-order pipeline and a register renaming stage, out-of-order issue/execution pipeline according to at least one embodiment of the invention. FIG. 4B is a block diagram illustrating an in-order architecture core and a register renaming logic, out-of-order issue/execution logic to be included in a processor according to at least one embodiment of the invention. The solid lined boxes in FIG. 4A illustrate the in-order pipeline, while the dashed lined boxes illustrates the register renaming, out-of-order issue/execution pipeline. Similarly, the solid lined boxes in FIG. 4B illustrate the in-order architecture logic, while the dashed lined boxes illustrates the register renaming logic and out-of-order issue/execution logic.

In FIG. 4A, a processor pipeline **400** includes a fetch stage **402**, a length decode stage **404**, a decode stage **406**, an allocation stage **408**, a renaming stage **410**, a scheduling (also known as a dispatch or issue) stage **412**, a register read/memory read stage **414**, an execute stage **416**, a write back/memory write stage **418**, an exception handling stage **422**, and a commit stage **424**.

In FIG. 4B, arrows denote a coupling between two or more units and the direction of the arrow indicates a direction of data flow between those units. FIG. 4B shows processor core **490** including a front end unit **430** coupled to an execution engine unit **450**, and both are coupled to a memory unit **470**.

The core **490** may be a reduced instruction set computing (RISC) core, a complex instruction set computing (CISC) core, a very long instruction word (VLIW) core, or a hybrid or alternative core type. As yet another option, the core **490** may be a special-purpose core, such as, for example, a network or communication core, compression engine, graphics core, or the like.

The front end unit **430** includes a branch prediction unit **432** coupled to an instruction cache unit **434**, which is coupled to an instruction translation lookaside buffer (TLB) **436**, which is coupled to an instruction fetch unit **438**, which is coupled to a decode unit **440**. The decode unit or decoder may decode instructions, and generate as an output one or

more micro-operations, micro-code entry points, microinstructions, other instructions, or other control signals, which are decoded from, or which otherwise reflect, or are derived from, the original instructions. The decoder may be implemented using various different mechanisms. Examples of suitable mechanisms include, but are not limited to, look-up tables, hardware implementations, programmable logic arrays (PLAs), microcode read only memories (ROMs), etc. The instruction cache unit **434** is further coupled to a level 2 (L2) cache unit **476** in the memory unit **470**. The decode unit **440** is coupled to a rename/allocator unit **452** in the execution engine unit **450**.

The execution engine unit **450** includes the rename/allocator unit **452** coupled to a retirement unit **454** and a set of one or more scheduler unit(s) **456**. The scheduler unit(s) **456** represents any number of different schedulers, including reservations stations, central instruction window, etc. The scheduler unit(s) **456** is coupled to the physical register file(s) unit(s) **458**. Each of the physical register file(s) units **458** represents one or more physical register files, different ones of which store one or more different data types, such as scalar integer, scalar floating point, packed integer, packed floating point, vector integer, vector floating point, etc., status (e.g., an instruction pointer that is the address of the next instruction to be executed), etc. The physical register file(s) unit(s) **458** is overlapped by the retirement unit **154** to illustrate various ways in which register renaming and out-of-order execution may be implemented (e.g., using a reorder buffer(s) and a retirement register file(s), using a future file(s), a history buffer(s), and a retirement register file(s); using a register maps and a pool of registers; etc.). Generally, the architectural registers are visible from the outside of the processor or from a programmer's perspective. The registers are not limited to any known particular type of circuit. Various different types of registers are suitable as long as they are capable of storing and providing data as described herein. Examples of suitable registers include, but are not limited to, dedicated physical registers, dynamically allocated physical registers using register renaming, combinations of dedicated and dynamically allocated physical registers, etc. The retirement unit **454** and the physical register file(s) unit(s) **458** are coupled to the execution cluster(s) **460**. The execution cluster(s) **460** includes a set of one or more execution units **462** and a set of one or more memory access units **464**. The execution units **462** may perform various operations (e.g., shifts, addition, subtraction, multiplication) and on various types of data (e.g., scalar floating point, packed integer, packed floating point, vector integer, vector floating point). While some embodiments may include a number of execution units dedicated to specific functions or sets of functions, other embodiments may include only one execution unit or multiple execution units that all perform all functions. The scheduler unit(s) **456**, physical register file(s) unit(s) **458**, and execution cluster(s) **460** are shown as being possibly plural because certain embodiments create separate pipelines for certain types of data/operations (e.g., a scalar integer pipeline, a scalar floating point/packed integer/packed floating point/vector integer/vector floating point pipeline, and/or a memory access pipeline that each have their own scheduler unit, physical register file(s) unit, and/or execution cluster—and in the case of a separate memory access pipeline, certain embodiments are implemented in which only the execution cluster of this pipeline has the memory access unit(s) **464**). It should also be understood that where separate pipelines are used, one or more of these pipelines may be out-of-order issue/execution and the rest in-order.

The set of memory access units **464** is coupled to the memory unit **470**, which includes a data TLB unit **472** coupled to a data cache unit **474** coupled to a level 2 (L2) cache unit **476**. In one exemplary embodiment, the memory access units **464** may include a load unit, a store address unit, and a store data unit, each of which is coupled to the data TLB unit **472** in the memory unit **470**. The L2 cache unit **476** is coupled to one or more other levels of cache and eventually to a main memory.

By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may implement the pipeline **400** as follows: 1) the instruction fetch **438** performs the fetch and length decoding stages **402** and **404**; 2) the decode unit **440** performs the decode stage **406**; 3) the rename/allocator unit **452** performs the allocation stage **408** and renaming stage **410**; 4) the scheduler unit(s) **456** performs the schedule stage **412**; 5) the physical register file(s) unit(s) **458** and the memory unit **470** perform the register read/memory read stage **414**; the execution cluster **460** perform the execute stage **416**; 6) the memory unit **470** and the physical register file(s) unit(s) **458** perform the write back/memory write stage **418**; 7) various units may be involved in the exception handling stage **422**; and 8) the retirement unit **454** and the physical register file(s) unit(s) **458** perform the commit stage **424**.

The core **490** may support one or more instructions sets (e.g., the x86 instruction set (with some extensions that have been added with newer versions); the MIPS instruction set of MIPS Technologies of Sunnyvale, Calif.; the ARM instruction set (with optional additional extensions such as NEON) of ARM Holdings of Sunnyvale, Calif.).

It should be understood that the core may support multithreading (executing two or more parallel sets of operations or threads), and may do so in a variety of ways including time sliced multithreading, simultaneous multithreading (where a single physical core provides a logical core for each of the threads that physical core is simultaneously multithreading), or a combination thereof (e.g., time sliced fetching and decoding and simultaneous multithreading thereafter such as in the Intel® Hyperthreading technology).

While register renaming is described in the context of out-of-order execution, it should be understood that register renaming may be used in an in-order architecture. While the illustrated embodiment of the processor also includes a separate instruction and data cache units **434/474** and a shared L2 cache unit **476**, alternative embodiments may have a single internal cache for both instructions and data, such as, for example, a Level 1 (L1) internal cache, or multiple levels of internal cache. In some embodiments, the system may include a combination of an internal cache and an external cache that is external to the core and/or the processor. Alternatively, all of the cache may be external to the core and/or the processor.

FIG. **5** is a block diagram of a single core processor and a multicore processor **500** with integrated memory controller and graphics according to embodiments of the invention. The solid lined boxes in FIG. **5** illustrate a processor **500** with a single core **502A**, a system agent **510**, a set of one or more bus controller units **516**, while the optional addition of the dashed lined boxes illustrates an alternative processor **500** with multiple cores **502A-N**, a set of one or more integrated memory controller unit(s) **514** in the system agent unit **510**, and an integrated graphics logic **508**.

The memory hierarchy includes one or more levels of cache within the cores, a set or one or more shared cache units **506**, and external memory (not shown) coupled to the set of integrated memory controller units **514**. The set of

shared cache units **506** may include one or more mid-level caches, such as level 2 (L2), level 3 (L3), level 4 (L4), or other levels of cache, a last level cache (LLC), and/or combinations thereof. While in one embodiment a ring based interconnect unit **512** interconnects the integrated graphics logic **508**, the set of shared cache units **506**, and the system agent unit **510**, alternative embodiments may use any number of well-known techniques for interconnecting such units.

In some embodiments, one or more of the cores **502A-N** are capable of multi-threading. The system agent **510** includes those components coordinating and operating cores **502A-N**. The system agent unit **510** may include for example a power control unit (PCU) and a display unit. The PCU may be or include logic and components needed for regulating the power state of the cores **502A-N** and the integrated graphics logic **508**. The display unit is for driving one or more externally connected displays.

The cores **502A-N** may be homogenous or heterogeneous in terms of architecture and/or instruction set. For example, some of the cores **502A-N** may be in order while others are out-of-order. As another example, two or more of the cores **502A-N** may be capable of execution the same instruction set, while others may be capable of executing only a subset of that instruction set or a different instruction set.

The processor may be a general-purpose processor, such as a Core™ i3, i5, i7, 2 Duo and Quad, Xeon™, Itanium™, XScale™ or StrongARM™ processor, which are available from Intel Corporation, of Santa Clara, Calif. Alternatively, the processor may be from another company, such as ARM Holdings, Ltd, MIPS, etc. The processor may be a special-purpose processor, such as, for example, a network or communication processor, compression engine, graphics processor, co-processor, embedded processor, or the like. The processor may be implemented on one or more chips. The processor **500** may be a part of and/or may be implemented on one or more substrates using any of a number of process technologies, such as, for example, BiCMOS, CMOS, or NMOS.

FIGS. **6-8** are exemplary systems suitable for including the processor **500**, while FIG. **9** is an exemplary system on a chip (SoC) that may include one or more of the cores **502**. Other system designs and configurations known in the arts for laptops, desktops, handheld PCs, personal digital assistants, engineering workstations, servers, network devices, network hubs, switches, embedded processors, digital signal processors (DSPs), graphics devices, video game devices, set-top boxes, micro controllers, cell phones, portable media players, hand held devices, and various other electronic devices, are also suitable. In general, a huge variety of systems or electronic devices capable of incorporating a processor and/or other execution logic as disclosed herein are generally suitable.

Referring now to FIG. **6**, shown is a block diagram of a system **600** in accordance with one embodiment of the present invention. The system **600** may include one or more processors **610, 615**, which are coupled to graphics memory controller hub (GMCH) **620**. The optional nature of additional processors **615** is denoted in FIG. **6** with broken lines.

Each processor **610,615** may be some version of the processor **500**. However, it should be noted that it is unlikely that integrated graphics logic and integrated memory control units would exist in the processors **610,615**. FIG. **6** illustrates that the GMCH **620** may be coupled to a memory **640** that may be, for example, a dynamic random access memory (DRAM). The DRAM may, for at least one embodiment, be associated with a non-volatile cache.

The GMCH **620** may be a chipset, or a portion of a chipset. The GMCH **620** may communicate with the processor(s) **610, 615** and control interaction between the processor(s) **610, 615** and memory **640**. The GMCH **620** may also act as an accelerated bus interface between the processor(s) **610, 615** and other elements of the system **600**. For at least one embodiment, the GMCH **620** communicates with the processor(s) **610, 615** via a multi-drop bus, such as a frontside bus (FSB) **695**.

Furthermore, GMCH **620** is coupled to a display **645** (such as a flat panel display). GMCH **620** may include an integrated graphics accelerator. GMCH **620** is further coupled to an input/output (I/O) controller hub (ICH) **650**, which may be used to couple various peripheral devices to system **600**. Shown for example in the embodiment of FIG. **6** is an external graphics device **660**, which may be a discrete graphics device coupled to ICH **650**, along with another peripheral device **670**.

Alternatively, additional or different processors may also be present in the system **600**. For example, additional processor(s) **615** may include additional processor(s) that are the same as processor **610**, additional processor(s) that are heterogeneous or asymmetric to processor **610**, accelerators (such as, e.g., graphics accelerators or digital signal processing (DSP) units), field programmable gate arrays, or any other processor. There can be a variety of differences between the physical resources **610, 615** in terms of a spectrum of metrics of merit including architectural, micro-architectural, thermal, power consumption characteristics, and the like. These differences may effectively manifest themselves as asymmetry and heterogeneity amongst the processors **610, 615**. For at least one embodiment, the various processors **610, 615** may reside in the same die package.

Referring now to FIG. **7**, shown is a block diagram of a second system **700** in accordance with an embodiment of the present invention. As shown in FIG. **7**, multiprocessor system **700** is a point-to-point interconnect system, and includes a first processor **770** and a second processor **780** coupled via a point-to-point interconnect **750**. Each of processors **770** and **780** may be some version of the processor **500** as one or more of the processors **610,615**.

While shown with only two processors **770, 780**, it is to be understood that the scope of the present invention is not so limited. In other embodiments, one or more additional processors may be present in a given processor.

Processors **770** and **780** are shown including integrated memory controller units **772** and **782**, respectively. Processor **770** also includes as part of its bus controller units point-to-point (P-P) interfaces **776** and **778**; similarly, second processor **780** includes P-P interfaces **786** and **788**. Processors **770, 780** may exchange information via a point-to-point (P-P) interface **750** using P-P interface circuits **778, 788**. As shown in FIG. **7**, IMCs **772** and **782** couple the processors to respective memories, namely a memory **732** and a memory **734**, which may be portions of main memory locally attached to the respective processors.

Processors **770, 780** may each exchange information with a chipset **790** via individual P-P interfaces **752, 754** using point to point interface circuits **776, 794, 786, 798**. Chipset **790** may also exchange information with a high-performance graphics circuit **738** via a high-performance graphics interface **739**.

A shared cache (not shown) may be included in either processor or outside of both processors, yet connected with the processors via P-P interconnect, such that either or both

processors' local cache information may be stored in the shared cache if a processor is placed into a low power mode.

Chipset **790** may be coupled to a first bus **716** via an interface **796**. In one embodiment, first bus **716** may be a Peripheral Component Interconnect (PCI) bus, or a bus such as a PCI Express bus or another third generation I/O interconnect bus, although the scope of the present invention is not so limited.

As shown in FIG. 7, various I/O devices **714** may be coupled to first bus **716**, along with a bus bridge **718** which couples first bus **716** to a second bus **720**. In one embodiment, second bus **720** may be a low pin count (LPC) bus. Various devices may be coupled to second bus **720** including, for example, a keyboard and/or mouse **722**, communication devices **727** and a storage unit **728** such as a disk drive or other mass storage device which may include instructions/code and data **730**, in one embodiment. Further, an audio I/O **724** may be coupled to second bus **720**. Note that other architectures are possible. For example, instead of the point-to-point architecture of FIG. 7, a system may implement a multi-drop bus or other such architecture.

Referring now to FIG. 8, shown is a block diagram of a third system **800** in accordance with an embodiment of the present invention. Like elements in FIGS. 7 and 8 bear like reference numerals, and certain aspects of FIG. 7 have been omitted from FIG. 8 in order to avoid obscuring other aspects of FIG. 8.

FIG. 8 illustrates that the processors **870**, **880** may include integrated memory and I/O control logic ("CL") **872** and **882**, respectively. For at least one embodiment, the CL **872**, **882** may include integrated memory controller units such as that described above in connection with FIGS. 5 and 7. In addition, CL **872**, **882** may also include I/O control logic. FIG. 8 illustrates that not only are the memories **832**, **834** coupled to the CL **872**, **882**, but also that I/O devices **814** are also coupled to the control logic **872**, **882**. Legacy I/O devices **815** are coupled to the chipset **890**.

Referring now to FIG. 9, shown is a block diagram of a SoC **900** in accordance with an embodiment of the present invention. Similar elements in FIG. 5 bear like reference numerals. Also, dashed lined boxes are optional features on more advanced SoCs. In FIG. 9, an interconnect unit(s) **902** is coupled to: an application processor **910** which includes a set of one or more cores **902A-N** and shared cache unit(s) **906**; a system agent unit **910**; a bus controller unit(s) **916**; an integrated memory controller unit(s) **914**; a set or one or more media processors **920** which may include integrated graphics logic **908**, an image processor **924** for providing still and/or video camera functionality, an audio processor **926** for providing hardware audio acceleration, and a video processor **928** for providing video encode/decode acceleration; an static random access memory (SRAM) unit **930**; a direct memory access (DMA) unit **932**; and a display unit **940** for coupling to one or more external displays.

FIG. 10 illustrates a processor containing a central processing unit (CPU) and a graphics processing unit (GPU), which may perform at least one instruction according to one embodiment. In one embodiment, an instruction to perform operations according to at least one embodiment could be performed by the CPU. In another embodiment, the instruction could be performed by the GPU. In still another embodiment, the instruction may be performed through a combination of operations performed by the GPU and the CPU. For example, in one embodiment, an instruction in accordance with one embodiment may be received and decoded for execution on the GPU. However, one or more operations within the decoded instruction may be performed

by a CPU and the result returned to the GPU for final retirement of the instruction. Conversely, in some embodiments, the CPU may act as the primary processor and the GPU as the co-processor.

In some embodiments, instructions that benefit from highly parallel, throughput processors may be performed by the GPU, while instructions that benefit from the performance of processors that benefit from deeply pipelined architectures may be performed by the CPU. For example, graphics, scientific applications, financial applications and other parallel workloads may benefit from the performance of the GPU and be executed accordingly, whereas more sequential applications, such as operating system kernel or application code may be better suited for the CPU.

In FIG. 10, processor **1000** includes a CPU **1005**, GPU **1010**, image processor **1015**, video processor **1020**, USB controller **1025**, UART controller **1030**, SPI/SDIO controller **1035**, display device **1040**, memory interface controller **1045**, MIPI controller **1050**, flash memory controller **1055**, dual data rate (DDR) controller **1060**, security engine **1065**, and I<sup>2</sup>S/I<sup>2</sup>C controller **1070**. Other logic and circuits may be included in the processor of FIG. 10, including more CPUs or GPUs and other peripheral interface controllers.

One or more aspects of at least one embodiment may be implemented by representative data stored on a machine-readable medium which represents various logic within the processor, which when read by a machine causes the machine to fabricate logic to perform the techniques described herein. Such representations, known as "IP cores" may be stored on a tangible, machine readable medium ("tape") and supplied to various customers or manufacturing facilities to load into the fabrication machines that actually make the logic or processor. For example, IP cores, such as the Cortex™ family of processors developed by ARM Holdings, Ltd. and Loongson IP cores developed the Institute of Computing Technology (ICT) of the Chinese Academy of Sciences may be licensed or sold to various customers or licensees, such as Texas Instruments, Qualcomm, Apple, or Samsung and implemented in processors produced by these customers or licensees.

FIG. 11 shows a block diagram illustrating the development of IP cores according to one embodiment. Storage **1130** includes simulation software **1120** and/or hardware or software model **1110**. In one embodiment, the data representing the IP core design can be provided to the storage **1130** via memory **1140** (e.g., hard disk), wired connection (e.g., internet) **1150** or wireless connection **1160**. The IP core information generated by the simulation tool and model can then be transmitted to a fabrication facility where it can be fabricated by a 3<sup>rd</sup> party to perform at least one instruction in accordance with at least one embodiment.

In some embodiments, one or more instructions may correspond to a first type or architecture (e.g., x86) and be translated or emulated on a processor of a different type or architecture (e.g., ARM). An instruction, according to one embodiment, may therefore be performed on any processor or processor type, including ARM, x86, MIPS, a GPU, or other processor type or architecture.

FIG. 12 illustrates how an instruction of a first type is emulated by a processor of a different type, according to one embodiment. In FIG. 12, program **1205** contains some instructions that may perform the same or substantially the same function as an instruction according to one embodiment. However the instructions of program **1205** may be of a type and/or format that is different or incompatible with processor **1215**, meaning the instructions of the type in program **1205** may not be able to executed natively by the

processor **1215**. However, with the help of emulation logic, **1210**, the instructions of program **1205** are translated into instructions that are natively capable of being executed by the processor **1215**. In one embodiment, the emulation logic is embodied in hardware. In another embodiment, the emulation logic is embodied in a tangible, machine-readable medium containing software to translate instructions of the type in the program **1205** into the type natively executable by the processor **1215**. In other embodiments, emulation logic is a combination of fixed-function or programmable hardware and a program stored on a tangible, machine-readable medium. In one embodiment, the processor contains the emulation logic, whereas in other embodiments, the emulation logic exists outside of the processor and is provided by a third party. In one embodiment, the processor is capable of loading the emulation logic embodied in a tangible, machine-readable medium containing software by executing microcode or firmware contained in or associated with the processor.

FIG. **13** is a block diagram contrasting the use of a software instruction converter to convert binary instructions in a source instruction set to binary instructions in a target instruction set according to embodiments of the invention. In the illustrated embodiment, the instruction converter is a software instruction converter, although alternatively the instruction converter may be implemented in software, firmware, hardware, or various combinations thereof. FIG. **13** shows a program in a high level language **1302** may be compiled using an x86 compiler **1304** to generate x86 binary code **1306** that may be natively executed by a processor with at least one x86 instruction set core **1316**. The processor with at least one x86 instruction set core **1316** represents any processor that can perform substantially the same functions as a Intel processor with at least one x86 instruction set core by compatibly executing or otherwise processing (1) a substantial portion of the instruction set of the Intel x86 instruction set core or (2) object code versions of applications or other software targeted to run on an Intel processor with at least one x86 instruction set core, in order to achieve substantially the same result as an Intel processor with at least one x86 instruction set core. The x86 compiler **1304** represents a compiler that is operable to generate x86 binary code **1306** (e.g., object code) that can, with or without additional linkage processing, be executed on the processor with at least one x86 instruction set core **1316**. Similarly, FIG. **13** shows the program in the high level language **1302** may be compiled using an alternative instruction set compiler **1308** to generate alternative instruction set binary code **1310** that may be natively executed by a processor without at least one x86 instruction set core **1314** (e.g., a processor with cores that execute the MIPS instruction set of MIPS Technologies of Sunnyvale, Calif. and/or that execute the ARM instruction set of ARM Holdings of Sunnyvale, Calif.). The instruction converter **1312** is used to convert the x86 binary code **1306** into code that may be natively executed by the processor without an x86 instruction set core **1314**. This converted code is not likely to be the same as the alternative instruction set binary code **1310** because an instruction converter capable of this is difficult to make; however, the converted code will accomplish the general operation and be made up of instructions from the alternative instruction set. Thus, the instruction converter **1312** represents software, firmware, hardware, or a combination thereof that, through emulation, simulation or any other process, allows a processor or other electronic device that does not have an x86 instruction set processor or core to execute the x86 binary code **1306**.

Embodiments of the present invention involve instruction(s) to efficiently monitor loop trip count. In an embodiment, loop trip count information of a loop may be stored in a dedicated hardware buffer. The average loop trip count of the loop may be calculated based on the stored loop trip count information. Based on the average trip count, loop optimizations may be applied or removed from the loop. In an embodiment, the stored loop trip count information includes an identifier identifying the loop, a total loop trip count of the loop, and an exit count of the loop. In an embodiment, the loop optimization(s) may be based on a determination that the average loop trip count is above a threshold. In an embodiment, loop optimization(s) may be removed based on a determination that the average loop trip count is below a threshold. In an embodiment, a section of code identified for optimization may contain an atomic region. An atomic region may be a region of the code where all the statements in that region have to be executed as a whole. The atomic region of code may be demarcated by a commit statement. In an embodiment, execution of at least a commit instruction in the loop may be skipped.

In an embodiment, the stored loop trip count information may be removed based on a determination that the average loop trip count is above a threshold or below a threshold. In an embodiment, if the hardware buffer overflows, the stored loop trip count information in the hardware buffer may be replaced based on a cache algorithm.

FIG. **14** illustrates optimization of program code **1410** in an embodiment. Program code **1410** may refer to compiler code, optimization code, application code, library code, or any other known formulation of code. In an exemplary embodiment, program code **1410** may be executed by a computer processor, such as binary code that is ready for execution, dynamically compiled for execution on processor, and/or dynamically optimized to execute on a processor. In addition, the insertion of code (operations, function calls, etc) and optimization of the code is performed through execution of program code, such as compiler and/or optimization code. As an example, optimization code is dynamically executed on a processor at runtime to optimize program code just before execution of the program code on a processor.

In an embodiment, the program code **1410** to be optimized may first be identified as program code which needs optimization. For example, a specific instruction or demarcation is utilized to indicate a section of code to be optimized or would likely benefit from optimization. As another option, a programmer provides hints regarding sections of the program code, which are utilized by optimization code to identify sections for optimization. In another embodiment, a region is identified/selected based on profiling information. For example, program code is profiled during execution by hardware, software executing on the processor, firmware, or a combination thereof. Here, profiling of the code generates hints or modifies original software hints, as well as provide direct identification of regions for optimizations. In addition, a section of code is potentially identified by certain attributes, such as a specific type, format, or order of the code. As a specific illustrative example, code including loops are targeted for potential optimization. For example, the program code **1410** includes a loop as seen at statement **1414** where the code branches back to statement **1412** if variable P is a particular boolean value. And profiling of the loops during execution determines which of the loops should be optimized. Also, if the loop includes specific code, such as loads and stores, that are to be optimized, then the region including such code is identified for optimization. As

can be seen from exemplary program code **1410**, the code includes loads and stores capable of being hoisted and sunk out of the loop to optimize loop execution.

In an embodiment, the section of code identified for optimization may contain an atomic region. An atomic region is a region of the code where all the statements in that region have to be executed as a whole. The atomic region of code may be demarcated by a commit statement **1412**. However, any known method of indicating a region is atomic may be utilized.

Code optimization such as dynamic binary optimization, used in hardware/software co-designed architectures, relies on atomic region support to solve many tough issues in the optimized code, such as memory order violation, imprecise exception, etc. However, the atomic region commit enforces strict sequential consistence which not only becomes a barrier to optimize instructions across the commit boundary, but also causes stalls since the commit needs to wait until all memory references are globally visible.

In an embodiment, program code **1410** identified for optimization may be converted to the program code **1420**, **1430**, **1440**, and **1450**. Program code **1420-1450** performs the same function as program code **1410**, and may be more efficient than program code **1410**. Program code **1420-1450** utilizes a conditional commit branch (BCC instruction) to optimize the code **1410**. In an embodiment, a conditional commit branch skips a commit operation if there is enough speculative resources available for program execution until the next conditional commit operation. In FIG. **14**, the conditional commit operation is shown in the inner loop **1430** (as BCC), indicating that the loop **1430** is executed without any commit operations until speculative resources available for program execution are exhausted. When speculative resources are exhausted, the execution exits the inner loop **1430**, commits **1440**, and repeats the inner loop **1430** after executing an outer loop branch **1450**. Restructuring program code **1410** to use conditional commits 1) reduces stalls by skipping unnecessary commits, and 2) increases performance since multiple iterations of loops can be executed in parallel, through for example, pipelining.

However, unless there are multiple iterations of the inner loop **1430** (i.e., a high inner loop trip count), and the conditional commit is skipped many times, issues which affect performance may arise. For example, if the original program code loop's **1410** trip count is  $N$ , the original program code loop **1410** takes  $B$  (e.g., 5) cycles, each commit takes  $C$  (e.g., 2) cycles, the pipelined kernel **1430** takes  $S$  (e.g., 3) cycles to execute, and the pipelined ramp-up code **1420** and outer loop branch **1450** take  $R$  (e.g., 18) cycles, then the original program code loop **1410** will take  $N*(B+C)=7N$  cycles and the restructured code with conditional commits (**1420-1450**) will take  $N*S+C+R=3N+20$  cycles. Thus, the restructured code with conditional commits (**1420-1450**) will be more efficient if  $7N > 3N+20$ , or  $N > 5$ . In other words, the restructured code with conditional commits (**1420-1450**) will have to skip at least 5 consecutive commits in order to achieve a performance improvement. The trip count of the inner loop **1430** may be low if the inner loop **1430** uses large amounts of speculative resources, thus forcing frequent commits **1440**, and consequently, the execution of the outer loop branch **1450**.

In an embodiment, the trip count of optimized code (for example, restructured code using conditional commits **1420-1450**) may be monitored for low trip counts to check whether some of the optimizations should be reverted (for example, reverted to the original program code loop **1410**). FIG. **15** illustrates exemplary optimized code to track loop

trip count in an embodiment. In an embodiment, the loop trip count may be stored in a register without utilizing dedicated hardware support. A register (RN) may be initialized **1521** to an initial value such as 0. Within the loop **1530**, the value in the register may be incremented, so that each time the loop **1530** is executed, the value in the register reflects the iteration number of the loop **1530**. Upon exiting the loop two memory update operations **1541** may be performed in the commit section of the code **1540**. The first memory operation may increment a counter in memory (mem[BK]) by the value in the register RN (i.e., the current loop trip count value). The second memory operation may increment a second counter in memory (mem[EX]) by 1 to keep track of the number of times of exit of the loop **1530**. In an embodiment, the outer loop branch **1550** may include an instruction to reset the register RN back to an initial value such as zero (not shown). To obtain the average loop trip count, the total loop trip count (value in mem[BK]) may be divided by the number of times of exit from the loop (value in mem[EX]). However, maintaining a separate register without utilizing hardware support may impact performance because of the use of the additional register RN, the use of one instruction slot **1531** within loop **1530**, and the two memory update operations **1541**.

FIG. **16** illustrates exemplary optimized code to track loop trip count in an embodiment. In an embodiment, hardware support may efficiently keep track of the loop trip count without the overheads pertaining to the code illustrated in FIG. **15**. A conditional commit instruction (as described above) with additional functionality may be added to the optimized program code. The conditional commit instruction with the added functionality (TC\_BCC **1631**) instructs the computer processor to store loop trip count information in a dedicated hardware buffer. The hardware buffer (trip count buffer) may store 1) an identifier identifying the loop (for example, loop **1630**), 2) the respective total loop trip count, and 3) the exit count from the respective loop. The average loop trip count can be calculated by dividing the total loop trip count by the exit count from the loop. In an embodiment, the identifier identifying the loop may be the branch instruction pointer (IP) address of the TC\_BCC instruction **1631**. In an embodiment, at particular time intervals the trip count buffer may be examined and the trip count information may be updated to the meta-data associated with the respective loops.

In an embodiment, the trip count buffer may store trip count information of multiple loops, and may be of a limited size. If the buffer overflows, the data in the buffer may be replaced based on a cache algorithm, including, Least Recently Used, Most Recently Used, Pseudo-Least Recently Used, Random Replacement, Segmented Least Recently Used, 2-Way Set Associative, Least Frequently Used, Adaptive Replacement Cache, Multi Queue Caching Algorithm and Re-Reference Interval Prediction.

In an embodiment, based on the detected loop trip count of a particular loop of optimized code (for example, loop **1630**), the computer processor may (for example, utilizing binary optimization software) convert the optimized code back to code without the loop optimizations, for example, the non-optimized program code **1410** (FIG. **14**). For example, the processor may determine the loop trip count of loop **1630** from the trip count buffer and determine whether the trip count is high enough to offset the overhead of the optimized code (for example, overhead related to the execution of ramp up code and outer loop branch code as discussed in relation to FIG. **14**), and if the trip count is too

low to offset the overhead, convert the optimized code **1620-1650** to non-optimized code such as the code **1410** (FIG. **14**).

In an embodiment, to efficiently utilize the space in the trip count buffer, the trip count buffer may only store entries for loops which have a loop count lower than a threshold. Thus, in an embodiment, if the processor examines the trip count buffer for the loop count of an optimized loop, and does not find an entry for that loop, the processor may conclude that the trip count for that loop is higher than the threshold, and consequently, determine that the loop should not be transformed back to a non-optimized loop.

FIG. **17** illustrates exemplary code without loop optimizations which reduces commit overhead in an embodiment. In an embodiment, program code without loop optimizations, such as code **1410** (FIG. **14**), may utilize a conditional commit instruction (TC\_BCC\_OR\_CMIT **1702**) to reduce commit overhead. The TC\_BCC\_OR\_CMIT instruction may cause the processor to perform a commit operation if the code execution exits loop **1700**. In other words, if the condition to exit the loop is false, the code will be executed again from the beginning of the loop **1702**. If the condition to exit the loop is true, the processor will perform a commit.

In a further embodiment, the functionality to update a trip count buffer (i.e., a dedicated hardware buffer to store trip count information) as discussed in relation to FIG. **16** may be integrated with the TC\_BCC\_OR\_CMIT instruction. For example, each time the TC\_BCC\_OR\_CMIT instruction is executed and the processor determines that loop **1700** has to be executed again, a total loop trip count for loop **1700** may be incremented in the trip count buffer. In addition, each time code execution exits loop **1700**, a exit loop counter associated with loop **1700** may be incremented in the trip count buffer. The average loop trip count for loop **1700** may be calculated based on the total trip count and the exit loop count associated with loop **1700**.

In an embodiment, based on the detected loop trip count of a particular loop without loop optimizations (for example, loop **1700**), the computer processor may convert the code without loop optimizations to optimized code, for example, the optimized code **1620-1650** (FIG. **16**). For example, the processor may determine the average loop trip count of loop **1700** from the trip count buffer and determine whether the trip count is high enough to offset overhead associated with execution of an optimized version of the code (for example, overhead related to the execution of ramp up code and outer loop branch code as discussed in relation to FIG. **14**), and if the trip count is high enough to offset the overhead, convert the code without loop optimizations to optimized code, such as code **1620-1650** (FIG. **16**).

In an embodiment, to efficiently utilize the space in the trip count buffer, the trip count buffer may only store entries for loops which have a loop count higher than a threshold. Thus, in an embodiment, if the processor examines the trip count buffer for the loop count of a particular non-optimized loop, and does not find an entry for that loop, the processor may conclude that the trip count for that loop is lower than the threshold, and consequently, determine that the loop should not be transformed to an optimized loop.

FIG. **18** illustrates a method for transforming code with loops in an embodiment. In an embodiment, the TC\_BC-C\_CMIT and TC\_BCC instruction may be integrated into code to efficiently execute code with loops. The computer processor may first identify code with one or more loops (box **1802**). For each loop, the processor may estimate whether the loop has a high trip count (box **1804**). If a high loop trip count is not estimated, the loop is not modified (box

**1816**), and the processor may wait for a period of time and repeat **1804**. If a high loop trip count is estimated at **1804**, the loop may be transformed to use the TC\_BCC\_CMIT instruction (box **1806**). When executing the loop with the TC\_BCC\_CMIT instruction, commit operations may be skipped when possible, and loop trip count information may be stored in a hardware loop trip count buffer. The loop trip count buffer may be checked periodically to determine whether the loop has a high trip count (box **1808**). If a high loop trip count is not detected, the loop is not modified (box **1816**). If a high loop trip count is detected, the loop is optimized, and is modified to use the TC\_BCC instruction (box **1820**). When executing the loop with the TC\_BCC instruction, commit operations may be skipped when possible, and loop trip count information may be stored in a hardware loop trip count buffer. The loop trip count buffer may be checked periodically to determine whether the loop has a low trip count (box **1812**). If a low trip count is not detected, the loop is not modified (box **1816**). If a low trip count is detected, the loop optimizations may be removed (box **1814**), and the loop may be transformed to utilize TC\_BCC\_CMIT (box **1806**), and **1806-1816** may be repeated.

In an embodiment, at **1812**, if the processor determines that the loop includes multiple loop optimizations and a low loop trip count is detected, the processor may first determine whether some of the loop optimizations can be retained with reasonable performance. If so, the processor may retain the profitable optimizations and remove only the optimizations which may degrade performance. The processor may execute **1814** and remove all optimizations if it determines that none of the optimizations are profitable.

A person having ordinary skill in the art will appreciate that high and low loop trip counts are relative, and may vary based on the type of code and the type of hardware/software executing the code. For example, the trip count threshold needed for efficiently pipelining a loop may be different from the trip count needed for redundancy elimination.

Embodiments of the mechanisms disclosed herein may be implemented in hardware, software, firmware, or a combination of such implementation approaches. Embodiments of the invention may be implemented as computer programs or program code executing on programmable systems comprising at least one processor, a storage system (including volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device.

Program code may be applied to input instructions to perform the functions described herein and generate output information. The output information may be applied to one or more output devices, in known fashion. For purposes of this application, a processing system includes any system that has a processor, such as, for example; a digital signal processor (DSP), a microcontroller, an application specific integrated circuit (ASIC), or a microprocessor.

The program code may be implemented in a high level procedural or object oriented programming language to communicate with a processing system. The program code may also be implemented in assembly or machine language, if desired. In fact, the mechanisms described herein are not limited in scope to any particular programming language. In any case, the language may be a compiled or interpreted language.

One or more aspects of at least one embodiment may be implemented by representative instructions stored on a machine-readable medium which represents various logic within the processor, which when read by a machine causes the machine to fabricate logic to perform the techniques

described herein. Such representations, known as “IP cores” may be stored on a tangible, machine readable medium and supplied to various customers or manufacturing facilities to load into the fabrication machines that actually make the logic or processor.

Such machine-readable storage media may include, without limitation, non-transitory, tangible arrangements of articles manufactured or formed by a machine or device, including storage media such as hard disks, any other type of disk including floppy disks, optical disks, compact disk read-only memories (CD-ROMs), compact disk rewritable’s (CD-RWs), and magneto-optical disks, semiconductor devices such as read-only memories (ROMs), random access memories (RAMs) such as dynamic random access memories (DRAMs), static random access memories (SRAMs), erasable programmable read-only memories (EPROMs), flash memories, electrically erasable programmable read-only memories (EEPROMs), magnetic or optical cards, or any other type of media suitable for storing electronic instructions.

Accordingly, embodiments of the invention also include non-transitory, tangible machine-readable media containing instructions or containing design data, such as Hardware Description Language (HDL), which defines structures, circuits, apparatuses, processors and/or system features described herein. Such embodiments may also be referred to as program products.

In some cases, an instruction converter may be used to convert an instruction from a source instruction set to a target instruction set. For example, the instruction converter may translate (e.g., using static binary translation, dynamic binary translation including dynamic compilation), morph, emulate, or otherwise convert an instruction to one or more other instructions to be processed by the core. The instruction converter may be implemented in software, hardware, firmware, or a combination thereof. The instruction converter may be on processor, off processor, or part on and part off processor.

Thus, techniques for performing one or more instructions according to at least one embodiment are disclosed. While certain exemplary embodiments have been described and shown in the accompanying drawings, it is to be understood that such embodiments are merely illustrative of and not restrictive on the broad invention, and that this invention not be limited to the specific constructions and arrangements shown and described, since various other modifications may occur to those ordinarily skilled in the art upon studying this disclosure. In an area of technology such as this, where growth is fast and further advancements are not easily foreseen, the disclosed embodiments may be readily modifiable in arrangement and detail as facilitated by enabling technological advancements without departing from the principles of the present disclosure or the scope of the accompanying claims.

What is claimed is:

1. A non-transitory machine-readable medium having stored thereon an instruction, which if performed by a machine causes the machine to perform a method comprising:

storing loop trip count information of a loop in a hardware buffer dedicated to store trip count information;  
calculating an average loop trip count of the loop based on the stored loop trip count information; and  
removing any loop optimizations from the loop based on the average loop trip count.

2. The machine-readable medium of claim 1, wherein the stored loop trip count information includes an identifier identifying the loop, a total loop trip count of the loop, and an exit count of the loop.

3. The machine-readable medium of claim 1, further comprising executing a commit instruction in the loop before removing the loop optimization, wherein removing the loop optimization comprises skipping execution of the commit instruction.

4. The machine-readable medium of claim 1, wherein the method further comprises:

removing the stored loop trip count information upon determining that the average loop trip count is one of:  
a) above a threshold and b) below a threshold.

5. The machine-readable medium of claim 1, further comprising applying the loop optimization upon determining that the average loop trip count is above a first threshold, wherein the loop optimization is subsequently removed upon determining that the average loop trip count is below a second threshold.

6. The machine-readable medium of claim 1, further comprising, upon overflow of the hardware buffer, replacing the stored loop trip count information based on a cache algorithm.

7. The machine-readable medium of claim 3, wherein the commit instruction demarcates an atomic section of code and skipping the commit instruction causes the atomic section of code to not execute.

8. A method comprising:

storing loop trip count information of a loop in a hardware buffer dedicated to store trip count information;  
calculating an average loop trip count of the loop based on the stored loop trip count information; and  
removing any loop optimizations from the loop based on the average loop trip count.

9. The method of claim 8, wherein the stored loop trip count information includes an identifier identifying the loop, a total loop trip count of the loop, and an exit count of the loop.

10. The method of claim 8, further comprising, executing a commit instruction in the loop before removing the loop optimization, wherein removing the loop optimization comprises skipping execution of the commit instruction.

11. The method of claim 8, further comprising:

removing the stored loop trip count information upon determining that the average loop trip count is one of:  
a) above a threshold and b) below a threshold.

12. The method of claim 8, further comprising applying the loop optimization upon determining that the average loop trip count is above a first threshold, wherein the loop optimization is subsequently removed upon determining that the average loop trip count is below a second threshold.

13. The method of claim 8, wherein the method further comprises:

upon overflow of the hardware buffer replacing the stored loop trip count information based on a cache algorithm.

14. The machine-readable medium of claim 10, wherein the commit instruction demarcates an atomic section of code and skipping the commit instruction causes the atomic section of code to not execute.

15. A system comprising:

a dedicated hardware buffer to store loop trip count information of a loop;  
a computer processor to calculate an average loop trip count of the loop based on the stored loop trip count information, and remove any loop optimizations from the loop based on the average loop trip count.



**16.** The system of claim **15**, wherein the stored loop trip count information includes an identifier identifying the loop, a total loop trip count of the loop, and an exit count of the loop.

**17.** A system comprising:

a processor for executing computer instructions, the computer instructions comprising:

storing loop trip count information of a loop in a hardware buffer dedicated to store trip count information;

calculating an average loop trip count of the loop based on the stored loop trip count information; and

removing any loop optimizations from the loop based on the average loop trip count.

**18.** The system of claim **17**, wherein the stored loop trip count information includes an identifier identifying the loop, a total loop trip count of the loop, and an exit count of the loop.

**19.** The system of claim **17**, wherein the computer instructions further comprise executing a commit instruction in the loop before removing the loop optimization, wherein removing the loop optimization comprises skipping execution of the commit instruction.

**20.** The system of claim **17**, wherein the computer instructions further comprise:

removing the stored loop trip count information upon determining that the average loop trip count is one of:  
a) above a threshold and b) below a threshold.

**21.** The system of claim **17**, wherein the computer instructions further comprise applying the loop optimization upon

determining that the average loop trip count is above a first threshold, wherein the loop optimization is subsequently removed upon determining that the average loop trip count is below a second threshold.

**22.** The system of claim **21**, wherein the loop optimization is removed upon determining that the average loop trip count is below a second threshold.

**23.** The system of claim **17**, wherein the computer instructions further comprise:

upon overflow of the hardware buffer replacing the stored loop trip count information based on a cache algorithm.

**24.** The machine-readable medium of claim **19**, wherein the commit instruction demarcates an atomic section of code and skipping the commit instruction causes the atomic section of code to not execute.

**25.** A non-transitory machine-readable medium having stored thereon data, which if performed by at least one machine, causes the at least one machine to fabricate at least one integrated circuit to perform a method comprising:

storing loop trip count information of a loop in a hardware buffer dedicated to store trip count information;

calculating an average loop trip count of the loop based on the stored loop trip count information; and

removing any loop optimizations from the loop based on the average loop trip count.

**26.** The machine-readable medium of claim **25**, wherein the stored loop trip count information includes an identifier identifying the loop, a total loop trip count of the loop, and an exit count of the loop.

\* \* \* \* \*